

Forensic DHCP Information Extraction from Home Routers

Memory Forensics on SOHO / Enduser Embedded Routing and Gateway Systems
University of Amsterdam

Tobias Fiebig (tobias.fiebig@os3.nl)

August 6, 2013

Abstract

This document explores the feasibility and admissibility of a so far unrecognized source for digital evidence. The extraction of a suspect's home router's main memory to obtain valuable evidence is proposed and implemented. This method aims at providing time lines of devices appearing in a home network, and therefore possibly in that home, along with their owner. The technique is implemented and tested on the standards of modern volatile memory forensics. The results show that the proposed method is viable and may be extended to a wide range of devices.

Keywords: Volatile Memory; Forensics; DHCP; Home Router;

1 Introduction

At least in the European Union, the consumer broadband Internet access has seen a huge growth in recent years. Undoubtedly this growth will continue within reasonable bounds [25].

To enable an end-user to use multiple devices behind his designated broadband connection, most Western Internet Service Providers (ISPs) equip their customers with dedicated devices that handle the broadband connection as well as the distribution of the connection to the local network¹.

This is usually accomplished by employing a local RFC1918 [21] addressing schema, which is managed by the provided device using the Dynamic Host Configuration Protocol (DHCP) [6]. Connectivity to the Internet is then established by performing Network Address Translation (NAT) [24] on the provided device.

¹Called Customer Premise Equipment (CPE) by ISPs

Common protocols for establishing the connections with local clients include IEEE 802.1(z) [19, 9] and IEEE 802.11(abgn) [32, 33, 35].

These devices are usually referenced as SOHO or End-user Routers, or, more fitting to their actual functionality of providing gateway services between the protocols used in the local network and the protocols used on the broadband link, as gateways. This document will further reference them just as *home routers*.

As home routers handle and control the local network traffic in their associated network segment, they may be a viable target during a forensic investigation. This document will explore these possibilities.

1.1 Related Work

Home routers usually employ either a MIPS [18] or ARM [10] based architecture [8] along with a dedicatedly tailored firmware. These devices are generally considered embedded devices.

Due to common challenges for embedded device development and necessities associated therewith, the Joint Test Action Group (JTAG) IEEE 1149.1 [16] was developed. Breeuwsma et. al. [4] investigated in 2006 how this port can be utilized for the forensic imaging of embedded devices and whether these ports are sufficiently accessible in available devices to allow for a common use of these during forensic investigations. They concluded that this method of access is a viable option for the acquisition of non volatile internal memory like for example the flash storage of a device [4].

In 2010, Roeloffs and Van Eijk [31] investigated the use of JTAG techniques to forensically extract the Random Access Memory (RAM) of TomTom GPS navigation systems [31]. Although they were unable to acquire large amounts of qualitative data, this is the first attempt of extracting RAM

from an embedded device via JTAG for forensic purposes known to the author of this document.

This spans the arc to memory forensics on home routers. While investigations on the extraction of forensic information from penetrated and abused home routers already have been conducted [28, 29] and a considerable amount of work has been put into the acquisition of network forensic information from home routers fully controlled by the investigator [15], a transfer of the techniques presented by Roeloffs and Van Eijk on TomTom systems to a general approach viable for home routers has not yet been conducted.

This research will attempt a first step into the direction of a universal forensic technique, allowing the utilization of volatile memory on home routers during a forensic investigation aimed at the owner of said device. To sufficiently scope the research at hand, the area of research will be limited to the investigation of the presence of specific devices within a Local Area Network (LAN) controlled by a single home router.

2 Hypothesis

Due to the known limitations of flash based storage [3] the author of this paper assumes that all relevant state information in regard to the previously mentioned RFC1918 [21] address propagation via DHCP [6] are kept in memory.

These state recordings, commonly called lease files, are necessary for the functionality of the DHCP. For each device associated with the network, the home router is responsible for they have to contain at least the Media Access Control (MAC) address of a client device as well as the time stamp of the last address assignment operation and the assigned Internet Protocol (IP) address [6].

Hence it is the author's hypothesis that from the information extracted from the main memory of a home router a timeline can be established, showing the association and disassociation of devices in a forensically viable manner.

Therefore this research will conduct a set of explorations and experiments allowing the verification of said hypothesis.

2.1 Contribution to the Field

If the proposed hypothesis can be validated, presence of a certain device at a certain place at a certain time can be established. In the light of recent developments in the mobile phone sector the presence of a specific mobile device in a network can be associated with the presence of the owner in the vicinity of the network devices.

Investigators may be able to utilize this information to indicate association between two persons, one being the owner of a certain identified device, the other one being the owner of the router device.

As this technique utilizes the JTAG port of a device, no cooperation from the legitimate owner of a home router is necessary.

If the approach succeeds, the door to further investigations of the memory contents of home router devices is opened.

3 Forensic Requirements

The restriction of conducting the proposed research in a forensically sound manner imposes various constraints on the methods that may be applied. While some are only relevant during a real criminal investigation and are composed mostly of procedures that have to be implemented during the handling of the evidence produced by and the data the evidence is produced of, some have to be incorporated into the method itself.

While these aspects will be briefly discussed here, the following section will cover how each of the constraints contemplated here has been addressed in the design process of the developed methodology.

If the requirements can not be implemented in the proposed method, this section will cover why this is not the case, and how a procedure that would have to be implemented during a real criminal investigation could satisfy the constraints nevertheless.

This summary of requirements is mostly based on the work of Vömel and Freiling [34], who recently published a detailed analysis of the constraints that have to be implemented during a volatile memory acquisition. Furthermore, an aspect of legal proceedings, the reproducibility of the extraction of evidence, will be covered.

3.1 Correctness and Completeness

The first constraint set for a forensically viable volatile memory image determined by Vömel and Freiling consists of correctness and completeness.

According to Vömel and Freiling, an acquired memory image may be considered complete and correct, if and only if all values set in the physical memory are conserved *as is*, that means unmodified, to the created image [34, p. 131], and if and only if it holds that all values that were present in the source physical memory if and only if they have been present in said source memory [34, p. 131].

Although an image may be partially correct but not complete, if it does not contain all values present in the physical memory, but those values it holds are correct, an image that is not fully correct can not be complete [34, p. 131].

3.2 Atomicity

The second constraint of Vömel and Freiling is the atomicity of a memory image. Following their definition [34, Definition 5, p.132] a memory image may be considered atomic, if and only if the created image represents the state of all processes operating on the physical memory at the same point in time [34, p.132].

More direct, this is only the case if all intrinsic processes of the surveyed system are simultaneously frozen and no further changes to the state of all memory regions occurs by those processes before the memory is read.

Hence, if the first N regions of memory have been read as they were at time point T , all subsequent memory regions have to be read in the state they have been during T , i.e. no modifications to the memory regions may have been performed by concurrently performed operations on the target [34, p.132].

3.3 Integrity

The third and last constraint introduced by Vömel and Freiling is the integrity of a memory image. A created memory image may therefore be considered to have integrity, if and only if the content of all memory regions is preserved in the same state relative to a fixed time point prior to the operations performed to recover the image [34, Definition 6, p.132].

Simply put, the memory acquisition process may not alter the memory of a target. If the memory of a target is altered, those sections altered by the acquisition process have to be recognized as altered by the acquisition utility and hence their integrity is tainted.

This also implies that a memory image may in fact provide a partial integrity.

3.4 Reproducibility

Digital Evidence used in a court of law has to have been produced in a reproducible manner. If information is extracted from something, may it be Deoxyribonucleic acid (DNA) or in this case an Information Technology system, the defendant has to be able to contract his own independent expert witness, who can re-evaluate the peoples claims about a piece of evidence.

Naturally, this is a severe issue in the case of information extracted from a systems memory, which according to Farmer and Venema ranges on their Order of Volatility (OOV) on the second most volatile rank with an expected life span of around 10 nanoseconds [7, p. 6].

While in general an image of the memory content is produced during the initial investigation, which may be permanently archived and used for the investigation, the reproducibility of the process is commonly limited.

In their 2006 paper, Sutherland et. al. approached this issue by recommending following a procedure proposed in the "*Directors and corporate advisors' guide to digital investigations and evidence*" [23], which recommends that the process of extraction is documented and observed by an independent eye-witness [26].

Such a witness may then - according to [23] - be interrogated in court, allowing the testing of the witnesses recollections of the evidence extraction, supporting the claim that the process was not tainted.

Sutherland et. al. compare this process with the interrogation of a traffic police officer during a court case handling a traffic violation observed by said officer [26, p. 67].

4 Method

As it was not possible to investigate the validity of the hypothesis on all home routers available in the




Figure 1: Equipped eJTAG header on a 1043ND.




Figure 2: A DLC5 Cable is used to connect a TP-Link 1043ND with a standard PC.

market due to the vast amount of different solutions from different vendors, a single device was chosen for a proof of concept implementation.

4.1 Investigated Device

The device that was chosen is the TP-Link 1043ND², a small IEEE 802.1z capable device, which also allows clients to connect via IEEE 802.1[abgn].

This device was chosen due to its easy availability in the laboratory. Furthermore, it uses a MIPS instruction set based Central Processing Unit (CPU). The Extended JTAG standard (eJTAG) Version 2.6 present in the MIPS architecture allows easy access to the system's main memory [18]. The specific central processing unit used in this model is the Atheros AR9132-BC1E which is a member of the AR71xx family.

4.2 Physical Interconnect

The eJTAG interface present on the TP-Link 1043ND was used to extract the system's running memory. To interface with the eJTAG port, the initially not accessible pin-out of the TP-Link 1043ND was equipped with a 2x7 pin row as shown in Figure 1.

This port was utilized to connect a standard PC with the 1043ND via a so called unbuffered Xilinx DLC5 Cable [36]. This cable is a fully passive connection, hence eliminates possible issues for the correctness of the obtained memory image as described in Section 3.1 due to the performance of

operations in the programmable logic of an more advanced cable.

A schema of the logical connections for that cable as described in [36] can be found in Appendix A.

The cable was then interfaced with the 1043ND and a standard PC's parallel port (LTP) as seen in Figure 2, which was used to perform the data recovery tasks.

4.3 Extraction Software

To interface with the 1043ND over the eJTAG port, a special JTAG software is necessary. Due to its overall good documentation, the author chose OpenOCD³ [20] for this project.

The used OpenOCD version, OpenOCD 0.7.0, was compiled only with support for a DLC5 parallel port interface.

In conjunction with the distribution supplied information and configuration file for AR71xx based chip-sets, OpenOCD can be utilized to extract the random access memory of a running device.

The stock OpenOCD does provide the possibility of extracting bulk chunks of memory. However, it does so by instructing the Memory Management Unit (MMU) of the CPU to deliver the desired memory contents. This is an issue for most of the constraints defined by Vömel and Freiling. Hence a method has to be implemented, that allows for the direct extraction of memory content from the device, without requiring active participation of any part of the device.

²<http://www.tp-link.com/en/products/details/?model=TL-WR1043ND>

³<http://openocd.sourceforge.net/>

As the eJTAG utilized for MIPS generally supports this [18], it should be possible to add this functionality to OpenOCD. In 2010 Timo Juhani Lindfors published a patch to an earlier version of OpenOCD⁴ that allows MMU bypassed memory access. This patch can also be found in Appendix B.

As a newer OpenOCD version was utilized for the development of the method at hand, the patch had to be adjusted. The adjusted patch can be found in Appendix C.

4.4 Extraction Process

To extract the memory of the 1043ND the first step is issuing a *halt* command to the CPU with OpenOCD. This command sent via the JTAG interface immediately halts the execution of all instructions on the CPU.

In the next step OpenOCD is instructed to extract the the memory of this device in MMU bypass mode. Based on the supplied boot messages found on a serial console attached to the 1043ND and information obtained from the U-Boot boot-loader source code⁵ the logical memory offset was established to be `0x80000000`. The size of the physical memory in the 1043ND is 32MB `0x2000000`, limited by the extends of the built-in chip.

With the applied physical dump patch OpenOCD will actually extract from `0xA0000000`, the physical memory location [27, p. 42ff], effectively bypassing the MMU.

Hence the command `dump_image phys img.bin 0x80000000 0x2000000` is send to OpenOCD. The memory image is extracted and then saved in the current working directory of OpenOCD with the filename supplied during the extraction process, in this case `img.bin`. The average speed during this process is around 0.65 KiB per second.

5 Validation

The proposed method has to hold up to the previously presented fundamentals of volatile memory forensic to be admissible during a forensic investigation. Hence it will be thoroughly investigated

if the method complies with the presented constraints, and if it does not fully comply with those constraints, in how far the possible issues reduce the admissibility.

5.1 Correctness

The correctness of the method is highly dependent on the correctness of the OpenOCD source code. Although the original intention of OpenOCD was not forensic soundness, the initial developer aimed at a tool that “[...] never displays wrong or inaccurate information” [20, p. 38].

Even if this claim is not evaluated within the thesis documenting the initial development of OpenOCD, the author of this document assumes that this requirement is fulfilled by OpenOCD, as it would diminish the usability of a debugger if it was unfulfilled. This assumption has to be made during this research, as the thorough investigation of the source code would exceed the scope of this project.

Another important aspect of correctness, not mentioned in the 2012 paper of Vömel and Freiling is the question whether what is found in the physical memory of a device is actually correct.

Extensive research on the impact of cosmic rays on integrated circuits has been performed in the last decades [5, 14, 37, 12].

While the flipping of single bits does not necessarily pose an issue in a normal volatile memory forensic investigation, the proposed method is highly focused on the correctness of a specific, small set of bits without additional parity information. Although events that may alter those bits are considerably sparse, they may become an issue during a forensic investigation.

Specifically, Ibe et. al. could establish in viable simulations that the process of down-scaling of chips increases the cosmic ray induced soft-error frequency. Their simulations found an increase of factor 6-7 for a migration step from 130nm to 22nm [12]. With the progressing decrease in chip size these issues hence may become more problematic.

Furthermore, it has been indicated in the literature that the altitude of a location has a considerable impact on the amount of observed soft errors due to an increase in the neutron flux [30].

Although these issues can not be denied, they are not investigated in-depth in the research at hand, as this would violate the intended scope of

⁴<http://lists.berlios.de/pipermail/openocd-development/2010-November/017278.html>

⁵<http://www.denx.de/wiki/U-Boot>

the project. Further research on this matter will be advised in a later section of this document.

5.2 Completeness

The memory image created from the running device is complete, given the method itself is also correct, if all bytes present in the physical memory are also written to the memory image.

The memory image is obtained starting from the previously known offset of the physical memory. The amount of bytes read exactly equals the storage capacity in the utilized physical memory chip⁶.

Hence, if the previously made assumptions hold, the proposed method is complete.

5.3 Atomicity

As previously described, the atomicity of an acquired memory snapshot is highly important [34]. Furthermore, the preserving of atomicity in the memory image acquisition process on full scale x86 based systems is a considerable challenge [22].

On the MIPS based system at hand however, this challenge is merely a minor issue. Although the extraction process is considerably slow - the extraction of one image of 32MB takes roughly 14 hours - the initial step of interrupting and freezing all processing performed on the CPU of the system ensures that no program running on that CPU is able to perform any operation on any memory area.

Therefore, even if atomicity is not ensured by the image creation process itself, but instead by a small action prior to the execution process, the atomicity of the obtained memory snapshot is preserved.

5.4 Integrity

The last constraint for a forensically viable volatile memory image is the integrity of said image. As previously described, a memory image has integrity, if and only if each of its subsets have the same state as exactly prior to the start of the acquisition process, i.e. if and only if the memory acquisition method does not alter the memory of the target system in its process.

Sadly OpenOCD requires a subset of the targets memory for its own processes if extended operations are performed on the system. Although according to the OpenOCD manual those operations do not include the simple memory extraction operations [1], it can not be finally be debarred, that the integrity of those memory areas has been tainted.

5.5 Reproducibility

As with most volatile memory forensics techniques, the general reproducibility of this method is highly limited. Procedural techniques as discussed in Section 3.4 may provide a sufficient documentation of the supplied method, so the reproducibility can be disregarded within the same parameters discussed for traditional volatile memory forensics methods.

Furthermore, if required in the specific case, more advanced methods could be implemented. Due to the small dimensions and the comparatively low power consumption of a device, it might be possible to conserve the device in halted state for extended time periods, if a constant power supply can be ensured. To implement such a power transfer, a utility similar to the one documented in US Patent 8,076,798 might be used [17].

5.6 Practical Image Verification

Besides the previously performed theoretical validation of the method, two other approaches for validating an obtained memory image have been tested.

The first approach is based on the work of Inoue, Adelstein and Joyce on self similarity in memory images and aims at the verification of the correctness of the proposed method. In 2011 Inoue, Adelstein and Joyce proposed the application of techniques known from biology to memory images [13]. They noticed that the presence of large amounts of self similarities within a memory image may yield that it is tainted. They researched the applicability of dotplots used in biology to illustrated self similarities to the analysis of memory images. To create such a dotplot, each page of a memory image⁷ is plotted against every other page of the image. The intersection is marked, if and only if those two pages are identical and the corresponding pages do not both consist of known initialization values of memory. Inoue, Adelstein

⁶For reference: WINBOND W9425G6JH-5

⁷4k on MIPS32

MAC Address	Hostname
52:54:00:40:f8:7b	death
52:54:00:c1:02:83	luggage
52:54:00:74:68:4d	poons
52:54:00:5c:af:ca	ridcully
52:54:00:6f:b0:32	rincewind
52:54:00:90:58:05	stibbons
52:54:00:00:ad:04	twoflower
52:54:00:d5:4e:33	vimmes

Table 1: List of used hosts with their associated MAC addresses.

and Joyce found that this method from biology is applicable to memory images and allows for the detection of tainted image acquisition techniques [13].

This technique proposed by Inoue, Adelstein and Joyce has been applied to memory images obtained with the method proposed by the author. One of the obtained dotplots can be seen in Figure 3. During the creation of the dotplot, all pages that contained only values related to memory initialization have been ignored. In this case these were pages exclusively containing either `0x00`, `0xFF`, `0x55` or `0xAA`.

The presented dotplot clearly shows no significant self similarity. This supports the conclusion of the theoretical evaluation that the proposed method is correct in a forensic context.

The second practical validation approach is related to the theoretical validation of the integrity of the proposed method. The higher the integrity of the method, the more similarity should exist between two images, subsequently taken from the same source.

Ideally both images should be identical bitwise.

The author performed those subsequent memory extractions two times. Both times the two images extracted within one process were bitwise identical.

5.7 Summary

After theoretically evaluating the previously postulated requirements for a forensically sound memory acquisition technique, the proposed method may be considered sufficiently viable in a forensic context.

A short practical evaluation further strengthens the theoretically obtained conclusions. Hence the




Figure 4: Schematic representation of the setup used for testing the proposed method.

author assumes the proposed method viable for a forensic memory extraction of volatile memory from home routers.

Therefore, the practical testing of the previously postulated hypothesis regarding the extractability of DHCP lease file information is the next step.

6 Empirical Verification

To empirically test the practical feasibility of the proposed method, the already presented TP-Link 1043ND was incorporated in an automated test setup. The setup would transparently simulate various scenarios and then extract the device memory with the proposed method. The extracted memory image can then be investigated.

As the simulated scenario is known prior to the investigation of the extracted memory images, the results of the investigation can be compared to the actual events, providing a metric of efficiency for this method.

6.1 Experiment Setup

The setup consisted of four main elements as shown in Figure 4. The core of the setup was the TP-Link 1043ND. A virtual machine server was connected to the internal network ports of the router, providing the virtual machines mimicking the network clients of the home router. A secondary system was connected to the JTAG port of the 1043ND and finally a Racktivity Remote Power Switch Unit⁸ provided the capability of remotely switching the 1043ND on and off, i.e. resetting it to the initial state.

⁸<http://www.racktivity.com/>




Figure 3: Dotplot showing self-similarity between pages in a memory image obtained by the author. The axis show the index of the corresponding pages.

The initial state in this case means that everything is powered off.

6.2 Data Acquisition

The experiments were then conducted using the Python script supplied in Appendix D. For each iteration of a simulated scenario, the whole setup would first be reset. This means that all virtual machines are stopped, and the home router was powered off for 120 seconds to erase all possible artifacts in memory that may survive a short cold or even hot reset of the device. The time-frame of 120 seconds was chosen, as research on cold boot attacks by Halderman et. al. suggests that the bits stored in modern SRAM chips decay to a state where non of the initial data is left after 50 seconds on standard operation temperature⁹ [11]. Hence the power-off time of 120 seconds should ensure that no patterns are left in the memory after a reset of the device.

After the initial setup reset, the 1043ND would be powered on. Based on the settings for each sce-

nario the virtual machines would then be powered on and off at will.

In total eight different scenarios were created for the purpose of this experiment. The timelines of the actual events can be found in Appendix E. These timelines are all relative to the initial power-on of the device. A rough outline of the conducted experiments can be found in Table 2

7 Results

During the investigation of the memory images from the 1043ND, it was discovered that the device seems to hold the lease file not on a memory filesystem as expected, but instead directly in the DHCP process' memory.

An example excerpt can be found in Appendix I. The first 48bit on line `0x01f691f0` e.g. show the beginning of the entry for one host, starting with the recognizable vendor id part of KVM hosts' MAC addresses, `52:54:00:XX:XX:XX`.

When compared with the MAC addresses used in the experiment as listed in Table 1, it becomes

⁹25.5°C to 44.1°C in the experiments of Halderman et. al.

Scenario	Description
adv-test-1-4	boot 1 host, shutdown, wait 4h, dump memory
adv-test-1-8	boot 1 host, shutdown, wait 8h, dump memory
adv-test-8-4	boot 8 hosts, shutdown, wait 4h, dump memory
adv-test-8-8	boot 8 hosts, shutdown, wait 8h, dump memory
plain-test-4	boot 4 hosts, dump memory
plain-test-8	boot 8 hosts, dump memory
complex	boot 3 hosts, wait 1.25h, boot 3 hosts, shutdown 2 hosts, wait 12h, dump memory

Table 2: Overview of the simulated scenarios.

even more obvious, it can be observed that the data found in the presented excerpt actually corresponds to the MAC addresses of hosts used in the experiments.

Although an extraction of the information would be possible, another angle of approach was found. The 1043ND firmware also logs the DHCP server related events to the system console. The content of said console can be found as plain ASCII values within the created memory dumps. An example of those loglines can be found in Appendix H.

As plain strings can be handled more easily, the syslog information was used for the timeline estimation.

Furthermore, issues with the global time correlation of the events taking place on the device were discovered. The model at hand seems to be unable to synchronize its local time with sources from the Internet. Hence no relation between real time and device relative time is easily possible.

Although tools like e.g. volatility¹⁰ support the extraction of the uptime, which could be used to establish a connection between the real and the device time, support for MIPS based memory dumps is not present.

To preserve the scope of this project, a different approach was taken, to estimate the overall device run time. As previously mentioned, the device system log information can be found in memory. To estimate the uptime, the oldest as well as the newest entry in the extracted syslog are used. The difference between the two points in time is then considered to be the uptime of the device. In conjunction with the known external time of the device, halt events with known device relative time can then be correlated to real times.

The extraction of information was conducted with a Python script designed specifically for this

purpose. This script can be found in Appendix G. The script creates visual and textual representations of the extracted timelines. Tables showing those timelines can be found in Appendix F.

From the created timelines, four different metrics were extracted.

1. Correctly detected clients.
2. Correctly detected join-time.
3. Amount of MACs in DHCPD memory.
4. Amount of false positives.

The detection rate of all clients is the count of all hosts that were used in a scenario and have been detected in the image.

The detection accuracy of device join-time is the amount of hosts where the detected network join-time corresponds to the host's boot time. Due to the boot process, a timeframe of 120 seconds is accepted as slack.

The amount of false positives consists of all hosts that have been determined to be present from the memory dump, although they have not been present, and the list of host in memory are all hosts for which MAC addresses have been found in the memory region suspected to hold the DHCP lease file. The latter operation has been performed manually.

7.1 Extracted Information

The two simple scenarios show a very good detection rate of the connected hosts. All hosts used in the experiment could also be identified in the associated memory dump using the developed tooling.

The four advanced scenarios show far more different results. While the presence of a single host with only one active host overall was reliably detected even after four and seven hours, the detection rate in the case of seven hosts being present was considerably low.

¹⁰<http://code.google.com/p/volatility/>

Scenario	Detection Rate	Accuracy	Hosts In-Memory	False Positives	Total Hosts
adv-test-1-4	1	1	1	0	1
adv-test-1-8	1	1	1	0	1
adv-test-8-4	2	2	8	0	8
adv-test-8-8	2	2	8	0	8
plain-test-4	4	4	4	0	4
plain-test-8	8	8	8	0	8
complex	6	3	6	0	6

Table 3: Results for the seven scenarios in the three different metrics.

This may however be related to the choice of the system log messages as data-source. A manual investigation of the corresponding memory images could produce the MAC addresses of all hosts utilized in this experiment from the memory. As expected, based on the previous observation all of them resided in coherent block of the memory dump, which is therefore assumed to belong to the DHCP server process on the home router.

The complex test, in which a more complicated set of events was created, could be identified in so far, that there was no occurrence of a host being determined to be present during a timeframe when it was in fact not present, i.e. no false positives occurred.

Furthermore, an estimate of when a host was present could be established.

To summarize the results, it can be held up that timelines could be created for various scenarios. Although some weaknesses exist on older data, these may be avoidable by utilizing more aspects of the memory resource.

8 Conclusion

The obtained results show that the proposed method is viable. The method did not - very importantly - show false positives. In all cases all relevant hosts could be identified, either with the implemented method or with easily possible extensions. Further more the data extracted from the memory dumps allows the creation of admissible timelines within reasonable bounds.

Hence the initially postulated hypothesis, that information extracted from the main memory of a home router should allow the creation of a timeline showing the association and disassociation of devices in a forensically viable manner, can be considered to be confirmed.

The proposed method is viable and should be

implementable for all available home routing devices that have an accessible eJTAG or JTAG port.

Although the base hypothesis has been confirmed, further research and engineering work is necessary to produce a market ready solution that allows the collection of forensically sound evidence from home routers.

8.1 Further Work

Many different chip-sets and firmware versions exist on the market. The investigation of these devices, how to interface with their specific JTAG ports - ideally in a soldering free manner¹¹ - and the creation of dedicated, tailored extraction utilities for all of these devices is probably the most important step.

This creation process also includes the investigation of different lease file formats and their extractability and informational value. As indicated in the results section, an additional incorporation of the actual lease file in the DHCP servers memory may have significantly increased the reliability and accuracy of the proposed method.

Furthermore, different angles of extractable information should be the matter of further research, as the overall impact of home router memory forensics could be highly increased by incorporating additional sources of information.

Home routing devices conduct various operations that may be of forensic interest and might allow a significant leap during the investigation of a case. Many devices do not only provide networking capabilities, but also further operations like Dynamic Name System (DNS) caching-resolver services or printing services.

These devices also handle all network traffic of a specific network segment. Due to the implemen-

¹¹A technique discussed on the Internet in 2009 may be viable here: <http://www.usbjtag.com/vbforum/showpost.php?p=18571>

tation specifics of NAT, the devices have to keep state for all connections performed via them [24]. In their 2011 paper, Beverly et. al. provided evidence that those network structures may be extracted from memory images, even hours after they actually happened [2].

An investigation like this could allow the determination of connections performed by devices that already have been shut down or removed from the vicinity before the investigation.

8.2 Defense Mechanisms

The matter of defending against this method is important, as this method might also be used by unlawful regimes against legitimate interests.

A defense against the presented methods is complicated. The DHCP lease information has to be stored somewhere. While dedicated systems with strict on-disk lease files and sophisticated measures against key extraction from memory via cold-boot attacks as described by Halderman et. al. [11] might be viable for a very small set of technologically advanced subjects, this does not help against this issue for most average customers.

A defense system for the masses could most probably only be implemented, if version 6 of the Internet Protocol, which allows a stateless address auto configuration, is used in conjunction with IPv6 privacy extensions, which masks the specific device identifier of a certain system.

That way no state has to be kept on a device, hence the target for possibly extractable information is highly reduced.

The most promising defense method is, however, the physical removal of JTAG access or disabling thereof in the chip.

Acknowledgments

Thomas Roth - For his support on Architectures.
Andreas Schuster - For feedback and supervision.

References

- [1] OpenOCD User's Guide 0.8.0-dev, Section 11.3, 2013.
<http://openocd.sourceforge.net/doc/html/CPU-Configuration.html>,
 accessed: Tue Jun 25 19:17:33 CEST 2013.
- [2] Robert Beverly, Simson Garfinkel, and Greg Cardwell. Forensic carving of network packets and associated data structures. *Digital Investigation*, 8:S78–S89, 2011.
- [3] Roberto Bez, Emilio Camerlenghi, Alberto Modelli, and Angelo Visconti. Introduction to flash memory. *Proceedings of the IEEE*, 91(4):489–502, 2003.
- [4] Ing Breeuwsma et al. Forensic imaging of embedded systems using JTAG (boundary-scan). *Digital Investigation*, 3(1):32–42, 2006.
- [5] V Degalahal, R Ramanarayanan, N Vijaykrishnan, Y Xie, and MJ Irwin. The effect of threshold voltages on the soft error rate [memory and logic circuits]. In *5th International Symposium on Quality Electronic Design*, 2004. *Proceedings.*, pages 503–508. IEEE, 2004.
- [6] R. Droms. Dynamic Host Configuration Protocol. RFC 2131 (Draft Standard), March 1997. Updated by RFCs 3396, 4361, 5494, 6842.
- [7] Dan Farmer and Wietse Venema. *Forensic Discovery*. Addison-Wesley Reading, 2005.
- [8] Robert Swope Fleming. *The end of the Intel age*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [9] Howard Frazier. The 802.3z gigabit Ethernet standard. *IEEE Network*, 12(3):6–7, 1998.
- [10] Steve B Furber. *ARM system Architecture*. Addison-Wesley Longman Publishing Co., Inc., 1996.
- [11] J Alex Halderman, Seth D Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A Calandrino, Ariel J Feldman, Jacob Appelbaum, and Edward W Felten. Let us remember: cold-boot attacks on encryption keys. *Communications of the ACM*, 52(5):91–98, 2009.
- [12] Eishi Ibe, Hitoshi Taniguchi, Yasuo Yahagi, Ki Shimbo, and Tadanobu Toba. Impact of scaling on neutron-induced soft error in SRAMs from a 250 nm to a 22 nm design rule. *IEEE Transactions on Electron Devices*, 57(7):1527–1538, 2010.
- [13] Hajime Inoue, Frank Adelstein, and Robert A Joyce. Visualization in testing a volatile memory forensic tool. *Digital Investigation*, 8:S42–S51, 2011.

- [14] Tanay Karnik and Peter Hazucha. Characterization of soft errors caused by single event upsets in CMOS processes. *IEEE Transactions on Dependable and Secure Computing*, 1(2):128–143, 2004.
- [15] Zhongli Liu, Yinjie Chen, Wei Yu, and Xinwen Fu. Generic network forensic data acquisition from household and small business wireless routers. In *2010 IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM)*, pages 1–6. IEEE, 2010.
- [16] Colin M Maunder and Rodham E Tulloss. *The test access port and boundary-scan architecture*. IEEE Computer Society Press Los Alamitos/Washington, DC, 1990.
- [17] Dean L Mehler and James P Wiebe. Uninterruptible a/c power supply transfer unit, December 13, 2011. US Patent 8,076,798.
- [18] MIPS Technologies Inc. *MIPS32 4KE™ Processor Core Family Software User’s Manual*. MIPS Technologies Inc., 2002.
- [19] Jon Postel and Joyce K Reynolds. Standard for the transmission of IP datagrams over IEEE 802 networks. 1988.
- [20] Dominic Rath. Open On-Chip Debugger, 2008.
- [21] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address Allocation for Private Internets. RFC 1918 (Best Current Practice), February 1996. Updated by RFC 6761.
- [22] Bradley Schatz. BodySnatcher: Towards reliable volatile memory acquisition by software. *Digital Investigation*, 4:126–134, 2007.
- [23] Peter Sommer. Directors and corporate advisors’ guide to Digital Investigations and evidence. 2005.
- [24] P. Srisuresh and M. Holdrege. IP Network Address Translator (NAT) Terminology and Considerations. RFC 2663 (Informational), August 1999.
- [25] Ryszard Struzak. Broadband internet in EU countries: limits to growth. *IEEE Communications Magazine*, 48(4):52–57, 2010.
- [26] Iain Sutherland, Jon Evans, Theodore Tryfonas, and Andrew Blyth. Acquiring volatile operating system data tools and techniques. *ACM SIGOPS Operating Systems Review*, 42(3):65–73, 2008.
- [27] Dominic Sweetman. *See MIPS run*. Morgan Kaufmann, 2010.
- [28] Patryk Szewczyk. ADSL Router Forensics Part 1: An introduction to a new source of electronic evidence. In *Australian Digital Forensics Conference*, page 13, 2007.
- [29] Patryk Szewczyk. ADSL Router Forensics Part 2: Acquiring Evidence. 2009.
- [30] Yoshiharu Tosaka, Ryoza Takasu, Taiki Uemura, Hideo Ehara, Hideya Matsuyama, Shigeo Satoh, Atsushi Kawai, and Masahiko Hayashi. Simultaneous measurement of soft error rate of 90 nm cmos sram and cosmic ray neutron spectra at the summit of mauna kea. In *IEEE International Reliability Physics Symposium, 2008. IRPS 2008.*, pages 727–728. IEEE, 2008.
- [31] Onno Van Eijk and Mark Roeloffs. Forensic acquisition and analysis of the Random Access Memory of TomTom GPS navigation systems. *Digital Investigation*, 6(3):179–188, 2010.
- [32] Richard Van Nee, Geert Awater, Masahiro Morikura, Hitoshi Takanashi, Mark Webster, and Karen W Halford. New high-rate wireless LAN standards. *IEEE Communications Magazine*, 37(12):82–88, 1999.
- [33] Dimitris Vassis, George Kormentzas, Angelos Rouskas, and Ilias Maglogiannis. The IEEE 802.11g standard for high data rate WLANs. *IEEE Network*, 19(3):21–26, 2005.
- [34] Stefan Vömel and Felix C Freiling. Correctness, atomicity, and integrity: Defining criteria for forensically-sound memory acquisition. *Digital Investigation*, 9(2):125–137, 2012.
- [35] Yang Xiao. IEEE 802.11 n: enhancements for higher throughput in wireless LANs. *IEEE Wireless Communications*, 12(6):82–91, 2005.
- [36] Xilinx Inc. JTAG Programmer Guide, Appendix B, 1999.
- [37] James F Ziegler, HW Curtis, HP Muhlfeld, CJ Montrose, B Chin, M Nicewicz, CA Russell, WY Wang, LB Freeman, P Hosier, et al. IBM experiments in soft fails in computer electronics (1978–1994). *IBM journal of research and development*, 40(1):3–18, 1996.

A DLC5 Xilinx Cable Simple Version

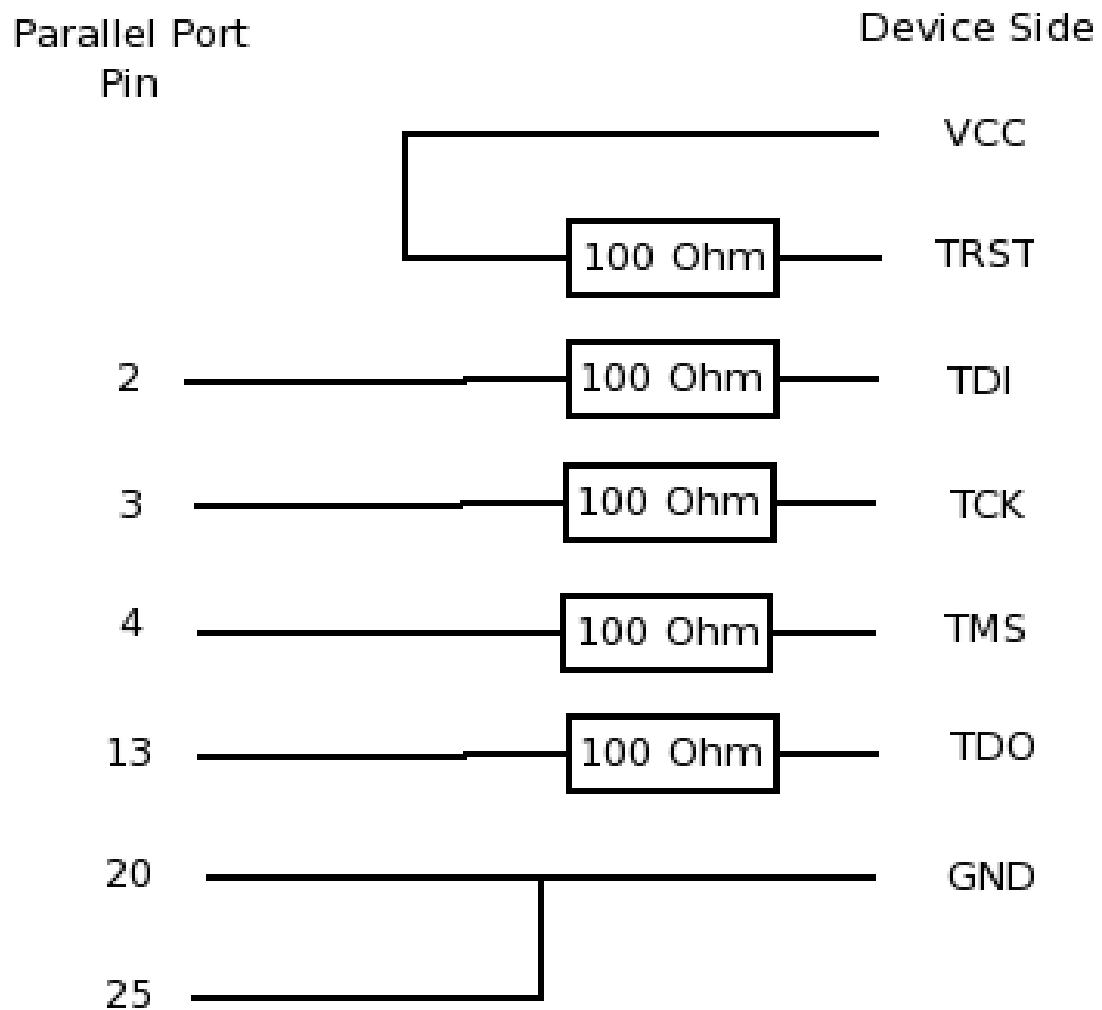


Figure 5: Schematics of the used simplified DLC5 Cable. The graphic is based on an image by the user "RealOpty" in the OpenWRT Wiki (<http://wiki.openwrt.org/doc/hardware/port.jtag.cable.unbuffered>) viewed on Mon Jul 1 21:15:41 CEST 2013.

B Unmodified Version of the OpenOCD DMA Patch

```
1 diff --git a/doc/openocd.texi b/doc/openocd.texi
2 index 70d789a..1b75f94 100644
3 --- a/doc/openocd.texi
4 +++ b/doc/openocd.texi
5 @@ -5731,9 +5731,12 @@ Otherwise, or if the optional @var{phys} flag is specified,
6   @cindex image dumping
7
8   @anchor{dump_image}
9   - at deffn Command {dump_image} filename address size
10  -Dump @var{size} bytes of target memory starting at @var{address} to the
11  -binary file named @var{filename}.
12  + at deffn Command {dump_image} [phys] filename address size
13  +Dump @var{size} bytes of target memory starting at @var{address} to
14  +the binary file named @var{filename}. When the current target has an
15  +MMU which is present and active, @var{addr} is interpreted as a
16  +virtual address. Otherwise, or if the optional @var{phys} flag is
17  +specified, @var{addr} is interpreted as a physical address.
18  @end deffn
19
20  @deffn Command {fast_load}
21  diff --git a/src/target/target.c b/src/target/target.c
22  index 93efa76..5cc1e6a 100644
23  --- a/src/target/target.c
24  +++ b/src/target/target.c
25  @@ -1393,12 +1393,22 @@ int target_write_buffer(struct target *target, uint32_t address,
26   uint32_t size,
27   * mode respectively, otherwise data is handled as quickly as
28   * possible
29   */
30  -int target_read_buffer(struct target *target, uint32_t address, uint32_t size, uint8_t
31  *buffer)
32  +static int target_read_buffer2(struct target *target, uint32_t address, uint32_t size,
33  uint8_t *buffer, bool physical)
34  {
35   int retval;
36   LOG_DEBUG("reading buffer of %i byte at 0x%8.8x",
37   (int)size, (unsigned)address);
38
39   + int (*read_fn)(struct target *target,
40   + uint32_t address, uint32_t size, uint32_t count, uint8_t *buffer);
41   + if (physical)
42   + {
43   +   read_fn=target_read_phys_memory;
44   + } else
45   + {
46   +   read_fn=target_read_memory;
47   + }
48   +
49   if (!target_was_examined(target))
50   {
51   LOG_ERROR("Target not examined yet");
52   }
53  @@ -1420,7 +1430,7 @@ int target_read_buffer(struct target *target, uint32_t address, uint32_t
54  size, u
55
56   if (((address % 2) == 0) && (size == 2))
57   {
58   - return target_read_memory(target, address, 2, 1, buffer);
59   + return read_fn(target, address, 2, 1, buffer);
60   }
61
62   /* handle unaligned head bytes */
63  @@ -1431,7 +1441,7 @@ int target_read_buffer(struct target *target, uint32_t address, uint32_t
64  size, u
65   if (unaligned > size)
66   unaligned = size;
67
68  - if ((retval = target_read_memory(target, address, 1, unaligned, buffer)) != ERROR_OK)
```

```

63 + if ((retval = read_fn(target, address, 1, unaligned, buffer)) != ERROR_OK)
64     return retval;
65
66     buffer += unaligned;
67 @@ -1444,7 +1454,7 @@ int target_read_buffer(struct target *target, uint32_t address, uint32_t
    size, u
68     {
69         int aligned = size - (size % 4);
70
71 - if ((retval = target_read_memory(target, address, 4, aligned / 4, buffer)) != ERROR_OK)
72 + if ((retval = read_fn(target, address, 4, aligned / 4, buffer)) != ERROR_OK)
73     return retval;
74
75     buffer += aligned;
76 @@ -1456,7 +1466,7 @@ int target_read_buffer(struct target *target, uint32_t address, uint32_t
    size, u
77     if(size >=2)
78     {
79         int aligned = size - (size%2);
80 - retval = target_read_memory(target, address, 2, aligned / 2, buffer);
81 + retval = read_fn(target, address, 2, aligned / 2, buffer);
82     if (retval != ERROR_OK)
83         return retval;
84
85 @@ -1467,13 +1477,18 @@ int target_read_buffer(struct target *target, uint32_t address,
    uint32_t size, u
86     /* handle tail writes of less than 4 bytes */
87     if (size > 0)
88     {
89 - if ((retval = target_read_memory(target, address, 1, size, buffer)) != ERROR_OK)
90 + if ((retval = read_fn(target, address, 1, size, buffer)) != ERROR_OK)
91     return retval;
92     }
93
94     return ERROR_OK;
95 }
96
97 +int target_read_buffer(struct target *target, uint32_t address, uint32_t size, uint8_t
    *buffer)
98 +{
99 + target_read_buffer2(target, address, size, buffer, false);
100 +}
101 +
102 int target_checksum_memory(struct target *target, uint32_t address, uint32_t size, uint32_t*
    crc)
103 {
104     uint8_t *buffer;
105 @@ -2605,6 +2620,12 @@ COMMAND_HANDLER(handle_dump_image_command)
106     struct duration bench;
107     struct target *target = get_current_target(CMD_CTX);
108
109 + bool physical=strcmp(CMD_ARGV[0], "phys")==0;
110 + if (physical)
111 + {
112 +     CMD_ARGC--;
113 +     CMD_ARGV++;
114 + }
115     if (CMD_ARGC != 3)
116         return ERROR_COMMAND_SYNTAX_ERROR;
117
118 @@ -2622,7 +2643,7 @@ COMMAND_HANDLER(handle_dump_image_command)
119     {
120         size_t size_written;
121         uint32_t this_run_size = (size > 560) ? 560 : size;
122 - retval = target_read_buffer(target, address, this_run_size, buffer);
123 + retval = target_read_buffer2(target, address, this_run_size, buffer, physical);
124     if (retval != ERROR_OK)
125     {
126         break;

```

```
127 @@ -5305,7 +5326,7 @@ static const struct command_registration target_exec_command_handlers[]
128 = {
129     .name = "dump_image",
130     .handler = handle_dump_image_command,
131     .mode = COMMAND_EXEC,
132     .usage = "filename address size",
133     },
134     {
135     .name = "verify_image",
136     --
137 1.7.2.3
```


C Modified Version of the OpenOCD DMA Patch

```
1 diff --git a/openocd-0.7.0/src/target/target.c b/openocd-0.7.0/src/target/target.c
2 index ed1a2cc..870d757 100644
3 --- a/openocd-0.7.0/src/target/target.c
4 +++ b/openocd-0.7.0/src/target/target.c
5 @@ -1758,11 +1758,22 @@ static int target_write_buffer_default(struct target *target, uint32_t
6     address,
7     * mode respectively, otherwise data is handled as quickly as
8     * possible
9     */
10 -int target_read_buffer(struct target *target, uint32_t address, uint32_t size, uint8_t
11 *buffer)
12 +static int target_read_buffer2(struct target *target, uint32_t address, uint32_t size,
13 uint8_t *buffer, bool physical)
14 {
15     LOG_DEBUG("reading buffer of %i byte at 0x%8.8x",
16               (int)size, (unsigned)address);
17 +
18 +// int (*read_fn)(struct target *target,
19 +//     uint32_t address, uint32_t size, uint32_t count, uint8_t *buffer);
20 + if (physical)
21 + {
22 +     read_fn=target_read_phys_memory;
23 + } else
24 + {
25 +     read_fn=target_read_memory;
26 + }
27 +
28 + if (!target_was_examined(target)) {
29     LOG_ERROR("Target not examined yet");
30     return ERROR_FAIL;
31 }
32 @@ -1787,7 +1798,7 @@ static int target_read_buffer_default(struct target *target, uint32_t
33 address, u
34 int retval = ERROR_OK;
35
36 if (((address % 2) == 0) && (size == 2))
37 - return target_read_memory(target, address, 2, 1, buffer);
38 + return read_fn(target, address, 2, 1, buffer);
39
40 /* handle unaligned head bytes */
41 if (address % 4) {
42 @@ -1796,7 +1807,7 @@ static int target_read_buffer_default(struct target *target, uint32_t
43 address, u
44 if (unaligned > size)
45     unaligned = size;
46
47 - retval = target_read_memory(target, address, 1, unaligned, buffer);
48 + retval = read_fn(target, address, 1, unaligned, buffer);
49 if (retval != ERROR_OK)
50     return retval;
51
52 @@ -1810,6 +1821,7 @@ static int target_read_buffer_default(struct target *target, uint32_t
53 address, u
54 int aligned = size - (size % 4);
55
56 retval = target_read_memory(target, address, 4, aligned / 4, buffer);
57 + retval = read_fn(target, address, 4, aligned / 4, buffer);
58 if (retval != ERROR_OK)
59     return retval;
60
61 @@ -1821,7 +1833,8 @@ static int target_read_buffer_default(struct target *target, uint32_t
62 address, u
63 /*prevent byte access when possible (avoid AHB access limitations in some cases)*/
64 if (size >= 2) {
65     int aligned = size - (size % 2);
66     - retval = target_read_memory(target, address, 2, aligned / 2, buffer);
67 + retval = read_fn(target, address, 2, aligned / 2, buffer);
```

```

61 +
62     if (retval != ERROR_OK)
63         return retval;
64
65 @@ -1831,7 +1844,7 @@ static int target_read_buffer_default(struct target *target, uint32_t
address, u
66     }
67     /* handle tail writes of less than 4 bytes */
68     if (size > 0) {
69 -     retval = target_read_memory(target, address, 1, size, buffer);
70 +     retval = read_fn(target, address, 1, size, buffer);
71     if (retval != ERROR_OK)
72         return retval;
73     }
74 @@ -1839,6 +1852,11 @@ static int target_read_buffer_default(struct target *target, uint32_t
address, u
75     return ERROR_OK;
76 }
77
78 +int target_read_buffer(struct target *target, uint32_t address, uint32_t size, uint8_t
*buffer)
79 +{
80 + return target_read_buffer2(target, address, size, buffer, false);
81 +}
82 +
83 int target_checksum_memory(struct target *target, uint32_t address, uint32_t size, uint32_t*
crc)
84 {
85     uint8_t *buffer;
86 @@ -2888,6 +2906,14 @@ COMMAND_HANDLER(handle_dump_image_command)
87     struct duration bench;
88     struct target *target = get_current_target(CMD_CTX);
89
90 + bool physical=strcmp(CMD_ARGV[0], "phys")==0;
91 + if (physical)
92 + {
93 +     CMD_ARGC--;
94 +     CMD_ARGV++;
95 + }
96 +
97 +
98     if (CMD_ARGC != 3)
99         return ERROR_COMMAND_SYNTAX_ERROR;
100
101 @@ -2910,7 +2936,7 @@ COMMAND_HANDLER(handle_dump_image_command)
102     while (size > 0) {
103         size_t size_written;
104         uint32_t this_run_size = (size > buf_size) ? buf_size : size;
105 -     retval = target_read_buffer(target, address, this_run_size, buffer);
106 +     retval = target_read_buffer2(target, address, this_run_size, buffer, physical);
107     if (retval != ERROR_OK)
108         break;
109
110 @@ -5650,7 +5676,7 @@ static const struct command_registration target_exec_command_handlers[]
= {
111     .name = "dump_image",
112     .handler = handle_dump_image_command,
113     .mode = COMMAND_EXEC,
114 -     .usage = "filename address size",
115 +     .usage = "['phys'] filename address size",
116 },
117 {
118     .name = "verify_image",
119 diff --git a/openocd-0.7.0/src/target/target.h b/openocd-0.7.0/src/target/target.h
120 index e6b931d..00d6925 100644
121 --- a/openocd-0.7.0/src/target/target.h
122 +++ b/openocd-0.7.0/src/target/target.h
123 @@ -587,4 +587,8 @@ void target_handle_event(struct target *t, enum target_event e);
124
125 extern bool get_target_reset_nag(void);

```

```
126 |
127 | /* adding definition for read_fn */
128 | int (*read_fn)(struct target *target, uint32_t address, uint32_t size, uint32_t count,
129 | +
130 | +
131 | #endif /* TARGET_H */
```

D Utility for Automated Method Tests

```
1 #!/usr/bin/env python
2 # -*- coding: utf-8 -*-
3
4 import warnings
5 warnings.filterwarnings("ignore")
6
7 import sys
8 import os
9 import re
10 import telnetlib
11 import libvirt
12 import time
13 import datetime
14 import subprocess
15 #import tempfile
16 #from optparse import OptionParser
17
18 libvirt_domains = ['ridcully', 'vimmes', 'death', 'rincewind', 'poons', 'luggage',
19 'twoflower', 'stibbons']
20 ram_offset = 0x80000000
21 ram_size = 0x2000000
22 #ram_size = 0x2000
23 debug = False
24
25 # global log directive
26 def log(str_f, line):
27     now = datetime.datetime.now()
28     f = open('./logs/'+str_f, "a+")
29     f.write(str(now) + ": " + line + "\n")
30     print str_f+": " + str(now) + ": " + line
31
32 def print_percent(str_file):
33     f_max = float(ram_size / 1024)
34     f_cur = 0
35     cur_pref = 0
36     stale_counter = 0
37
38     print "processing "+str_file+" : "
39     while f_cur < f_max:
40         proc = subprocess.Popen(["du ~/ocd/"+str_file], stdout=subprocess.PIPE, shell=True)
41         (out, err) = proc.communicate()
42
43         f_cur = float(out.strip().split()[0])
44
45         f_per = (f_cur/f_max * 100)
46
47         sys.stdout.flush()
48         sys.stdout.write("[ "+"#"*int(f_per * 0.78)+" "*int((100-int(f_per-1))*0.78) +" ]
49         "+str(round(f_per,2))+"%\r")
50         sys.stdout.flush()
51         time.sleep(1)
52         if cur_pref == f_cur:
53             stale_counter = stale_counter + 1
54         elif cur_pref < f_cur:
55             stale_counter = 0
56         if stale_counter > 60:
57             f_cur = f_max
58     print ""
59
60 def print_percent_sleep(s_time):
61     f_max = float(s_time)
62     f_cur = 0.0
63     while f_cur <= f_max:
64         f_per = (f_cur/f_max * 100)
65         sys.stdout.flush()
66         sys.stdout.write("[ "+"#"*int(f_per * 0.78)+" "*int((100-int(f_per-1))*0.78) +" ]   cur:
67         "+str(f_cur)+"s (" +str(round(f_per,2))+"%\r")
```

```

65     sys.stdout.flush()
66     time.sleep(1)
67     f_cur = f_cur + 1
68     print ""
69
70
71 # racktivity controll section
72 def set_port(str_host, str_port, str_user, str_pass, str_unit, int_port, int_state):
73     ret = "Port "+str(int_port)+" on "+str_unit+" set to "+str(int_state)+": "
74     tn = telnetlib.Telnet(str_host, str_port)
75     tn.read_until("Login: ")
76     tn.write(str_user + "\n")
77     tn.read_until("Password: ")
78     tn.write(str_pass + "\n")
79
80     tn.read_until("Login successful.")
81
82     tn.write("\n")
83     tn.read_until("Login: ")
84     tn.write(str_user + "\n")
85     tn.read_until("Password: ")
86     tn.write(str_pass + "\n")
87
88     tn.read_until("PROMPT>")
89
90
91     tn.write("SET "+str(str_unit)+" PORTSTAT "+str(int_port)+" "+str(int_state)+"\n")
92     tn.read_until("\n")
93     ret += tn.read_until("\n").strip()
94     tn.read_until("PROMPT>")
95     tn.write("LOGOUT\n")
96
97     tn.read_until("Command OK:")
98     tn.close()
99     return ret
100
101 def port_on(int_port):
102     return set_port('192.168.23.2', '2001', 'USERNAME', 'PASSWORD', 'P1', int_port, 1)
103
104 def port_off(int_port):
105     return set_port('192.168.23.2', '2001', 'USERNAME', 'PASSWORD', 'P1', int_port, 0)
106
107 # memory aquisition
108 def get_mem_generic(str_host, str_port, str_file_prfx, hex_offset, hex_size):
109     ret = ""
110
111     str_file =
112     str_file_prfx+"-"+str(hex(hex_offset)).strip('L')+"-"+str(hex(hex_size)).strip('L')+".bin"
113
114     tn = telnetlib.Telnet(str_host, str_port)
115     tn.read_until(">")
116     tn.write("reset\n")
117     tn.read_until(">")
118     tn.write("halt\n")
119     tn.read_until(">")
120     tn.write("dump_image phys "+str_file+" "+str(hex(hex_offset)).strip('L')+"
121     "+str(hex(hex_size)).strip('L')+"\n")
122     print_percent(str_file)
123     tn.read_until("\n")
124     ret = tn.read_until("\n")
125     tn.read_until(">")
126     tn.write("exit\n")
127     tn.close()
128
129     return ret.strip() + " to: "+str_file
130
131 def get_mem(str_file_prfx, hex_offset, hex_size):
132     return get_mem_generic('127.0.0.1', '4444', str_file_prfx, hex_offset, hex_size)
133
134 # libvirt controls

```

```

133 def libvirt_test():
134     conn=libvirt.open("qemu+tcp://192.168.23.1/system")
135     print conn.numOfDomains()
136     print conn.listDefinedDomains()
137     print conn.listDomainsID()
138     conn.close()
139
140 def lv_start_n(number):
141     conn=libvirt.open("qemu+tcp://192.168.23.1/system")
142     vms = conn.listDefinedDomains()
143     ret = []
144     for idx in range(0,number):
145         d = conn.lookupByName(vms[idx])
146         ret.append(vms[idx])
147         d.create()
148     return ret
149
150 def lv_stop_all():
151     conn=libvirt.open("qemu+tcp://192.168.23.1/system")
152     ret = "Stopped "+str(conn.numOfDomains())+" domains."
153     ids = conn.listDomainsID()
154     for dom in ids:
155         d = conn.lookupByID(dom)
156         d.destroy()
157     conn.close()
158     return ret
159
160 def lv_stop_host(host):
161     conn=libvirt.open("qemu+tcp://192.168.23.1/system")
162     ret = "Stopped "+host
163     d = conn.lookupByName(host)
164     d.destroy()
165     conn.close()
166     return ret
167
168 # utility functions
169 def reset():
170     ll = "reseting setup"
171     log("generic.txt", ll)
172     ll = port_off(1)
173     log("generic.txt", ll)
174     ll = lv_stop_all()
175     log("generic.txt", ll)
176     if not debug:
177         log("generic.txt", "Sleeping for 120s")
178         print_percent_sleep(120)
179     else:
180         log("generic.txt", "Sleeping for 10s")
181         print_percent_sleep(10)
182
183 def test_plain(int_hosts):
184     ll = port_on(1)
185     log("plain-test-"+str(int_hosts)+"-hosts.txt", ll)
186
187     if not debug:
188         hosts = lv_start_n(int_hosts)
189         ll = "started: "+str(hosts)
190         log("plain-test-"+str(int_hosts)+"-hosts.txt", ll)
191
192         log("plain-test-"+str(int_hosts)+"-hosts.txt", "Sleeping for 460s")
193         print_percent_sleep(460)
194     else:
195         log("plain-test-"+str(int_hosts)+"-hosts.txt", "Sleeping for 20s")
196         print_percent_sleep(20)
197
198     ll = get_mem("plain-test-"+str(int_hosts)+"-hosts", ram_offset, ram_size)
199     log("plain-test-"+str(int_hosts)+"-hosts.txt", ll)
200
201 def test_adv(int_hosts, int_dist):
202     ll = port_on(1)

```

```

203 log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", ll)
204
205 if not debug:
206     hosts = lv_start_n(int_hosts)
207     ll = "started: "+str(hosts)
208     log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", ll)
209
210     log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", "Sleeping for 460s")
211     print_percent_sleep(460)
212
213     ll = lv_stop_all()
214     log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", ll)
215
216     sleep = int_dist * 3600
217     log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", "Sleeping for
218         "+str(sleep)+"s")
219     print_percent_sleep(sleep)
220 else:
221     log("plain-test-"+str(int_hosts)+"-hosts.txt", "Sleeping for 20s")
222     print_percent_sleep(20)
223
224 ll = get_mem("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts", ram_offset, ram_size)
225 log("adv-test-"+str(int_hosts)+"-"+str(int_dist)+"-hosts.txt", ll)
226
227 def test_complex():
228     ll = port_on(1)
229     log("complex-test.txt", ll)
230
231     hosts = lv_start_n(3)
232     ll = "started: "+str(hosts)
233     log("complex-test.txt", ll)
234
235     log("complex-test.txt", "Sleeping for 4600s")
236     print_percent_sleep(4600)
237
238     hosts2 = lv_start_n(3)
239     ll = "started: "+str(hosts2)
240     log("complex-test.txt", ll)
241
242     ll = lv_stop_host(hosts[1])
243     log("complex-test.txt", ll)
244
245     ll = lv_stop_host(hosts[2])
246     log("complex-test.txt", ll)
247
248     log("complex-test.txt", "Sleeping for 43200s")
249     print_percent_sleep(43200)
250
251     ll = get_mem("complex-test", ram_offset, ram_size)
252     log("complex-test.txt", ll)
253
254 ## main method of experiment
255 def main():
256     # reset everything. Shutdown all vms, unpower the device.
257     reset()
258
259     for cnt in range(1,9):
260         test_plain(cnt)
261         reset()
262
263     for cnt in range(1,9):
264         for cnt2 in range(1,9):
265             test_adv(cnt,cnt2)
266             reset()
267
268     if not debug:
269         test_complex()
270         reset()
271
272 main()

```

E Real Event Timelines Extracted from Logs

E.1 adv-test-1-4

Time	52:54:00:5C:AF:CA
1	up
462	down
14879	memory dump

Table 4: Experiment Setting: One host is powered up and runs for approximately 460 seconds. Four hours later a memory dump is created.

E.2 adv-test-1-8

Time	52:54:00:5C:AF:CA
0	up
462	down
29293	memory dump

Table 5: Experiment Setting: One host is powered up and runs for approximately 460 seconds. Eight hours later a memory dump is created.

E.3 plain-test-4

Time	52:54:00:40:f8:7B	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:D5:4E:33
2	up	up	up	up
463	memory dump			

Table 6: Experiment Setting: Four hosts are booted. Approximately eight minutes later a memory dump is created.

E.4 adv-test-8-4

Time	52:54:00:00:AD:04	52:54:00:40:f8:7B	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:74:68:4D	52:54:00:90:58:05	52:54:00:C1:02:83	52:54:00:D5:4E:33
5 469	up down	up down	up down	up down	up down	up down	up down	up down
14884	memory dump							

Table 7: Experiment Setting: Eight hosts are powered up and runs for approximately 460 seconds. Four hours later a memory dump is created.

E.5 adv-test-8-8

Time	52:54:00:00:AD:04	52:54:00:40:f8:7B	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:74:68:4D	52:54:00:90:58:05	52:54:00:C1:02:83	52:54:00:D5:4E:33
5 468	up down	up down	up down	up down	up down	up down	up down	up down
29299	memory dump							

Table 8: Experiment Setting: Eight hosts are powered up and runs for approximately 460 seconds. Eight hours later a memory dump is created.

E.6 complex

Time	52:54:00:40:f8:7B	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:74:68:4D	52:54:00:C1:02:83	52:54:00:D5:4E:33
1	up	up				up
4609			up	up	up	down
4610	down					
47855	memory dump					

Table 9: Experiment Setting: Three hosts are initially powered up. Approximately 1.25 hours later two of them are powered down and three new hosts are booted.

E.7 plain-test-8

Time	52:54:00:00:AD:04	52:54:00:40:f8:7B	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:74:68:4D	52:54:00:90:58:05	52:54:00:C1:02:83	52:54:00:D5:4E:33
5	up	up	up	up	up	up	up	up
466	memory dump							

Table 10: Experiment Setting: Eight hosts are booted. Approximately eight minutes later a memory dump is created.

F Event Timelines as Extracted from Memory dumps

F.1 adv-test-1-4

Time	52:54:00:5C:AF:CA
45	DISCOVER
46	REQUEST
14883	memory dump

Table 11: Experiment Setting: One host is powered up and runs for approximately 460 seconds. Four hours later a memory dump is created.

F.2 adv-test-1-8

Time	52:54:00:5C:AF:CA
52	DISCOVER
53	REQUEST
29283	memory dump

Table 12: Experiment Setting: One host is powered up and runs for approximately 460 seconds. Eight hours later a memory dump is created.

F.3 plain-test-4

Time	52:54:00:D5:4E:33	52:54:00:5C:AF:CA	52:54:00:6F:B0:32	52:54:00:40:F8:7B
39	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST
40				
43				
44				
45				
46				
482	memory dump			

Table 13: Experiment Setting: Four hosts are booted. Approximately eight minutes later a memory dump is created.

F.4 adv-test-8-4

Time	52:54:00:5C:AF:CA	52:54:00:6F:B0:32
52	DISCOVER REQUEST	DISCOVER
53		
63		
14883	memory dump	

Table 14: Experiment Setting: Eight hosts are powered up and runs for approximately 460 seconds. Four hours later a memory dump is created.

F.5 adv-test-8-8

Time	52:54:00:D5:4E:33	52:54:00:5C:AF:CA
51	DISCOVER	DISCOVER REQUEST
52		
56		
29282	memory dump	

Table 15: Experiment Setting: Eight hosts are powered up and runs for approximately 460 seconds. Eight hours later a memory dump is created.

F.6 complex

Time	52:54:00:6F:B0:32	52:54:00:40:F8:7B	52:54:00:C1:02:83	52:54:00:D5:4E:33	52:54:00:74:68:4D	52:54:00:5C:AF:CA
47	REQUEST	DISCOVER	REQUEST	DISCOVER	REQUEST	REQUEST
48						
49		REQUEST				
38567						
38669	REQUEST	REQUEST	REQUEST	REQUEST	REQUEST	
38735						
40143						
41528						
41743	REQUEST	REQUEST	REQUEST	REQUEST	REQUEST	
42243						
43646						
44512						
45463	REQUEST	REQUEST	REQUEST	REQUEST	REQUEST	
45083						
46785						
47439						
47882	memory dump					

Table 16: Experiment Setting: Three hosts are initially powered up. Approximately 1.25 hours later two of them are powered down and three new hosts are booted.

F.7 plain-test-8

Time	52:54:00:00:AD:04	52:54:00:6F:B0:32	52:54:00:40:F8:7B	52:54:00:90:58:05	52:54:00:C1:02:83	52:54:00:D5:4E:33	52:54:00:74:68:4D	52:54:00:5C:AF:CA
54	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST	DISCOVER REQUEST		DISCOVER REQUEST
55								
63								
64								
65								
66								
67								
68								
69								
70							REQUEST	
483	memory dump							

Table 17: Experiment Setting: Eight hosts are booted. Approximately eight minutes later a memory dump is created.

G Utility for Lease Information Extraction

```
1 #!/usr/bin/env python
2 # -*- coding: utf-8 -*-
3
4 import warnings
5 warnings.filterwarnings("ignore")
6
7 import sys
8 import os
9 import re
10
11 import time
12 import datetime
13 import subprocess
14
15 import svgwrite
16 from svgwrite import cm, mm, rgb, deg
17
18 def parse_date(str_date):
19     re_non_space = re.compile(r'[^ ]+')
20     lst_date = re_non_space.findall(str_date)
21
22     dict_offsets = {'Jan':0, 'Feb':2678400, 'Mar':5097600, 'Apr':7776000, 'May':10368000,
23                    'Jun':13046400, 'Jul':15638400, 'Aug':18316800, 'Sep':20995200, 'Oct':23587200,
24                    'Nov':26265600, 'Dec':28857600}
25     lst_hms = lst_date[2].split(':')
26
27     return int(dict_offsets[lst_date[0]]) + 24 * 60 * 60 * (int(lst_date[1]) - 1) +
28           int(lst_hms[0]) * 60 * 60 + int(lst_hms[1]) * 60 + int(lst_hms[2])
29
30 def get_strings(str_file):
31     proc = subprocess.Popen(["strings "+str_file], stdout=subprocess.PIPE, shell=True)
32     (out, err) = proc.communicate()
33     return out
34
35 def get_dates(str_mem):
36     dict_ret = {}
37     re_mem = re.compile(r'^([JFMASOND][aepuco][nbrylgptvc].*)\ \ (none)\ \ kern.notice\
38 <25>DHCPS.*(DISCOVER\ from\
39 ([0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}
40 [:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2})|REQUEST\ from\
41 ([0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}
42 [:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}[:-][0-9a-fA-F]{2}))')
43     for str_line in str_mem.split('\n'):
44         if re_mem.match(str_line):
45             lst_data = re_mem.findall(str_line)
46             str_date = lst_data[0][0]
47             date = parse_date(str_date)
48             str_type = ""
49             str_obj = ""
50             if lst_data[0][3]:
51                 str_type = "REQUEST"
52                 str_obj = lst_data[0][3]
53             else:
54                 str_type = "DISCOVER"
55                 str_obj = lst_data[0][2]
56
57             dict_data = {"date":date, "type":str_type, "object":str_obj}
58             if dict_ret.has_key(date):
59                 dict_ret[date].append(dict_data)
60             else:
61                 dict_ret[date] = [dict_data]
62     return dict_ret
63
64 def create_timeline(str_file):
65     str_mem = get_strings(str_file)
```

```

60 dict_date = get_dates(str_mem)
61
62 lst_keys = dict_date.keys()
63 lst_keys.sort()
64
65 int_max_time = get_max_time(str_mem)
66
67 create_svg(dict_date, str_file.split('/')[7],int_max_time)
68
69
70 def get_uniq_macs(dict_times):
71     dict_macs = {}
72     lst_macs = []
73     counter = 0
74     for key in dict_times.keys():
75         for item in dict_times[key]:
76             lst_macs.append(item['object'])
77     lst_macs = list(set(lst_macs))
78     lst_macs.sort()
79
80     for mac in lst_macs:
81         dict_macs[mac] = counter
82         counter = counter + 1
83     return dict_macs
84
85 def get_max_time(str_mem):
86     max_time = 0
87     re_mem = re.compile(r'^([JFMASOND][aepuco][nbrylgptvc].*)\ \ (none)\ \ kern.notice')
88     for str_line in str_mem.split('\n'):
89         if re_mem.match(str_line):
90             lst_data = re_mem.findall(str_line)
91             str_date = lst_data[0]
92             date = parse_date(str_date)
93             if max_time < date:
94                 max_time = date
95     return max_time + 100
96
97 def scale_int(in_min, in_max, out_min, out_max, to_scl):
98     range_in = float(in_max - in_min)
99     range_out = float(out_max - out_min)
100
101     factor = range_out / range_in
102     return int( ((to_scl - in_min) * range_out) / range_in + out_min )
103
104
105 def create_svg(dict_times, str_image, int_max_time):
106     dict_macs = get_uniq_macs(dict_times)
107
108     print str_image
109
110     dwg = svgwrite.Drawing(filename='/tmp/svg/'+str_image.split('.')[0]+'.svg',
111                             size=(1280,(10+len(dict_macs.keys()))*10))
112     shapes = dwg.add(dwg.g(id='shapes', fill='red'))
113
114     # add base labels
115     for mac in dict_macs.keys():
116         shapes.add(dwg.text(mac, insert=(10 , 10 * int(dict_macs[mac]) + 10 ), fill='black',
117                               class_='text',font_size=9 ))
118
119     # add timeline
120     shapes.add(dwg.rect(insert=(100,10 * len(dict_macs.keys())), size=(1000,1), fill='black'))
121     shapes.add(dwg.text("Rel. Time t since device boot ( in seconds )", insert=(550, 80 + 10 *
122     len(dict_macs.keys())), fill='black', class_='text',font_size=9 ))
123     shapes.add(dwg.text("Black/Gray: Handed out leases, darkness determines certainty of
124     presence Red: Requests Only", insert=(550, 90 + 10 * len(dict_macs.keys())), fill='black',
125     class_='text',font_size=9 ))
126     for i in range(0,1500,500):
127         shapes.add(dwg.rect(insert=(100 + i ,10 * len(dict_macs.keys())), size=(2,10),
128                               fill='black'))
129     int_tmp_counter = 0

```



```

124     for j in str(scale_int(0, 1000, 0, int_max_time, i)):
125         shapes.add(dwg.text(j, insert=(99 + i , 20 + 10 * len(dict_macs.keys()) +
126             int_tmp_counter * 10), fill='black', class_='text',font_size=9 ))
127         int_tmp_counter = int_tmp_counter + 1
128     for i in range(0,1250,250):
129         shapes.add(dwg.rect(insert=(100 + i ,10 * len(dict_macs.keys())), size=(1,8),
130             fill='black'))
131         int_tmp_counter = 0
132         for j in str(scale_int(0, 1000, 0, int_max_time, i)):
133             shapes.add(dwg.text(j, insert=(99 + i , 20 + 10 * len(dict_macs.keys()) +
134                 int_tmp_counter * 10), fill='black', class_='text',font_size=9 ))
135             int_tmp_counter = int_tmp_counter + 1
136     for i in range(0,1050,50):
137         shapes.add(dwg.rect(insert=(100 + i ,10 * len(dict_macs.keys())), size=(1,4),
138             fill='black'))
139         int_tmp_counter = 0
140         for j in str(scale_int(0, 1000, 0, int_max_time, i)):
141             shapes.add(dwg.text(j, insert=(99 + i , 20 + 10 * len(dict_macs.keys()) +
142                 int_tmp_counter * 10), fill='black', class_='text',font_size=9 ))
143             int_tmp_counter = int_tmp_counter + 1
144     str_prt = "Time "
145     for mac in dict_macs.keys():
146         str_prt = str_prt + " & " + mac
147     print str_prt + "\\\\"
148     for timeslot in dict_times.keys():
149         for item in dict_times[timeslot]:
150             if item['type'] == "DISCOVER":
151                 offset = 100 + scale_int(0, int_max_time, 0, 1000, timeslot)
152                 hour = 4
153                 mac = item['object']
154                 if offset + hour > 1100:
155                     hour = hour - (offset + hour - 1100)
156                 shapes.add(dwg.rect(insert=(offset , 10 * int(dict_macs[mac])), size=(hour,9),
157                     fill='red'))
158     for item in dict_times[timeslot]:
159         if item['type'] == "REQUEST":
160             offset = 100 + scale_int(0, int_max_time, 0, 1000, timeslot)
161             hour = scale_int(0, int_max_time, 0, 1000, 3600)
162             mac = item['object']
163             if offset + hour > 1100:
164                 hour = hour - (offset + hour - 1100)
165             shapes.add(dwg.rect(insert=(offset , 10 * int(dict_macs[mac])), size=(4,9),
166                 fill='black'))
167             shapes.add(dwg.rect(insert=(offset , 10 * int(dict_macs[mac])), size=(hour,4),
168                 fill='black'))
169             offset = offset + hour
170             if offset + hour > 1100:
171                 hour = hour - (offset + hour - 1100)
172             shapes.add(dwg.rect(insert=(offset , 10 * int(dict_macs[mac])), size=(hour,4),
173                 fill='grey'))
174             offset = offset + hour
175             if offset + hour > 1100:
176                 hour = hour - (offset + hour - 1100)
177             shapes.add(dwg.rect(insert=(offset , 10 * int(dict_macs[mac])), size=(hour,4),
178                 fill='lightgrey'))
179     str_prt = str(item['date']) + " & "
180     #print timeslot
181     #print dict_times[timeslot]
182     for mac in dict_macs.keys():
183         last = ""
184         for item in dict_times[timeslot]:
185             if item['object'] == mac and not last == item['object']+item['type']:
186                 last = item['object']+item['type']
187                 str_prt = str_prt + item['type']
188             str_prt = str_prt + " & "
189     print str_prt[0:-2] + "\\\\"
190     print "\hline"

```

```
183 | print str(int_max_time)+" & \\multicolumn{" +str(len(dict_macs.keys()))+"}{c|}{memory dump}  
    | \\\\"  
184 | dwg.save()  
185 |  
186 | print ""  
187 |  
188 |  
189 | if __name__ == "__main__":  
190 |     create_timeline(sys.argv[1])
```

H Example System Log Messages from Memory Dump

```

1 Jan  1 13:06:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
2 Jan  1 13:06:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
3 Jan  1 13:10:39 (none) kern.notice <25>DHCP:Recv REQUEST from 52:54:00:6F:B0:32 86
4 Jan  1 13:10:39 (none) kern.notice <25>DHCP:Send ACK to 192.168.1.103 76
5 Jan  1 13:11:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
6 Jan  1 13:11:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
7 Jan  1 13:11:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
8 Jan  1 13:11:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
9 Jan  1 13:16:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
10 Jan  1 13:16:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
11 Jan  1 13:16:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
12 Jan  1 13:16:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
13 Jan  1 10:41:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
14 Jan  1 10:42:47 (none) kern.notice <25>DHCP:Recv REQUEST from 52:54:00:6F:B0:32 86
15 Jan  1 00:00:47 (none) kern.notice <25>DHCP:Recv DISCOVER from 52:54:00:40:F8:7B 87
16 Jan  1 00:00:48 (none) kern.notice <25>DHCP:Send OFFER with ip 192.168.1.105 83
17 Jan  1 00:00:48 (none) kern.notice <25>DHCP:Recv DISCOVER from 52:54:00:D5:4E:33 87
18 Jan  1 00:00:49 (none) kern.notice <25>DHCP:Send OFFER with ip 192.168.1.101 83
19 Jan  1 00:00:49 (none) kern.notice <25>DHCP:Recv REQUEST from 52:54:00:40:F8:7B 86
20 Jan  1 00:00:49 (none) kern.notice <25>DHCP:Send ACK to 192.168.1.105 76
21 Jan  1 00:00:49 (none) kern.notice Jan  1 00:00:49
22 Jan  1 12:01:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
23 Jan  1 12:06:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
24 Jan  1 12:06:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
25 Jan  1 12:06:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
26 Jan  1 12:06:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
27 Jan  1 12:07:26 (none) kern.notice <25>DHCP:Recv REQUEST from 52:54:00:74:68:4D 86
28 Jan  1 12:07:26 (none) kern.notice <25>DHCP:Send ACK to 192.168.1.104 76
29 Jan  1 12:11:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
30 Jan  1 12:11:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
31 Jan  1 12:11:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
32 Jan  1 12:11:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
33 Jan  1 12:16:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
34 Jan  1 12:16:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
35 Jan  1 12:16:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
36 Jan  1 12:16:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
37 Jan  1 12:21:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
38 Jan  1 12:21:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112
39 Jan  1 12:21:22 (none) kern.notice <25>DHCP:GET ip:9164665a mask:ffffffe0 gateway:91646641
  dns1:9164600b dns2:91646016 static route:0 140
40 Jan  1 12:21:22 (none) kern.notice <25>Dynamic IP(DHCP Client) obtained an IP successfully 96
41 Jan  1 12:21:52 (none) kern.notice <25>DHCP:Recv REQUEST from 52:54:00:6F:B0:32 86
42 Jan  1 12:21:52 (none) kern.notice <25>DHCP:Send ACK to 192.168.1.103 76
43 Jan  1 12:26:21 (none) kern.notice <25>DHCP:Send REQUEST to server 9164602a with request ip
  9164665a 107
44 Jan  1 12:26:22 (none) kern.notice <25>DHCP:Recv ACK from server 9164602a with ip 9164665a
  lease time 600 112

```

I Possible DHCPD in-memory Lease file

1	01d4f4c0	00 00 00 01 52 54 00 90	58 05 00 00 00 00 00 00RT..X.....
2	01d4f540	00 00 00 01 52 54 00 00	ad 04 00 00 00 00 00 00RT.....
3	01d4f5c0	00 00 00 01 52 54 00 c1	02 83 00 00 00 00 00 00RT.....
4	01d4f640	00 00 00 01 52 54 00 74	68 4d 00 00 00 00 00 00RT.thM.....
5	01d4f6c0	00 00 00 01 52 54 00 40	f8 7b 00 00 00 00 00 00RT.@.{.....
6	01d4f740	00 00 00 01 52 54 00 d5	4e 33 00 00 00 00 00 00RT..N3.....
7	01d4f940	00 00 00 01 52 54 00 5c	af ca 00 00 00 00 00 00RT.\.....
8	01d4f9c0	00 00 00 01 52 54 00 6f	b0 32 00 00 00 00 00 00RT.o.2.....
9	01e05af0	81 d9 52 54 00 00 00 00	00 00 00 00 81 e0 58 94	..RT.....X..
10	01edbf80	5f 50 4f 52 54 00 00 00	00 00 00 00 00 00 00 00	_PORT.....
11	01f50f90	00 00 00 06 52 54 00 90	58 05 00 00 00 00 00 08RT..X.....
12	01f50fd0	00 00 00 00 00 00 00 00	52 54 00 90 58 05 00 00RT..X...
13	01f58b30	50 4f 52 54 00 00 00 00	32 32 37 20 00 00 00 00	PORT....227
14	01f58b40	45 50 52 54 00 00 00 00	32 32 39 20 00 00 00 00	EPRT....229
15	01f5c8e0	81 e0 52 54 00 00 00 00	00 00 00 00 81 15 59 00	..RT.....Y..
16	01f66430	00 00 00 00 52 54 00 90	58 05 00 00 00 00 00 00RT..X.....
17	01f66910	00 00 00 00 52 54 00 90	58 05 00 00 00 00 00 00RT..X.....
18	01f691b0	00 00 00 00 52 54 00 5c	af ca 00 00 00 00 00 00RT.\.....
19	01f691f0	52 54 00 6f b0 32 00 00	00 00 00 00 00 00 00 00	RT.o.2.....
20	01f69260	00 00 00 00 00 00 00 00	52 54 00 40 f8 7b 00 00RT.@.{...
21	01f692a0	00 00 00 00 52 54 00 74	68 4d 00 00 00 00 00 00RT.thM.....
22	01f692e0	52 54 00 c1 02 83 00 00	00 00 00 00 00 00 00 00	RT.....
23	01f69350	00 00 00 00 00 00 00 00	52 54 00 90 58 05 00 00RT..X...