

# Gibraltar

## A Mobile Host-Based Intrusion Protection System

Grant A. Jacoby, Thadeus Hickman, Stuart P. Warders,  
Barak Griffin, Aaron Darensburg, Daniel E. Castle

Department of Electrical Engineering and Computer Science,  
United States Military Academy, West Point, NY 10996  
Email: dg6742@usma.edu; Phone: 1-845-938-6605; Fax 1-845-938-3807

Keys – Host-Based, Mobile, Security, Forensic, IDS

### Abstract

*Although mobile devices are globally omnipresent, security developments for these devices have not kept pace with their technological advancements. Thus, mobile devices are increasingly vulnerable to intrusions and malicious attacks. Gibraltar combats these growing threats by monitoring demands placed on battery current (mA) as well as correlating power and event activities, such as processes, open ports, and registry keys. This combination serves as an early warning tripwire-like sensor for mobile hosts, blocking as well as identifying attacks. The end state for Gibraltar is to provide a totally host-based proactive form of intrusion detection systems (IDS) that can be easily integrated into current network IDS to provide an enhancement in detecting, alerting and responding to various intrusions. This paper outlines the design, test, and build methodologies used to resolve attack-sensing and warning problems and discusses lessons learned.*

### I. Introduction

There is a need for an efficient host-centric method of intrusion detection in mobile wireless computing devices. With the increasing use of mobile devices in business and for personal information, these devices will be a likely target for attack. Moreover, most of these devices are by default poorly configured for security.

The purpose of Gibraltar is to design and implement a resource non-intensive hardware and software solution to combat the threat of attacks on mobile devices using their smart batteries as well as other features available in the device. Present day security measures make a large impact on a mobile device's resources. Gibraltar attempts to turn what many deem as a mobile device's greatest constraint, power consumption, into a strength. To accomplish this transformation, Gibraltar utilizes the battery signal as a reference to device usage

without the need to conduct more complex scans that persistently use processing cycles. Furthermore, Gibraltar integrates this host-based method of system defense with that of a distributed intrusion protection system (DIPS).

### II. Related Issues

Despite the fact that security threats against mobile devices are on the rise [1], security utilities are not widely configured into mobile devices by default. Most users have limited means of combating malicious activity. For example, current intrusion detection methods used for PCs were not designed with mobile platform in mind. Providing guarantees for sufficient processing performance and battery life on mobile devices are at a premium and mobile developers need to keep these constraints always in mind. Therefore, security development for mobile battery operated devices must take a different approach in combating malicious activity to maximize their security as well as their performance.

Gibraltar's approach to security considers the constraints of the mobile device. Conventional intrusion detection systems are ill suited due to the limited processing ability and power constraints of mobile devices. As the next two sections explain, Gibraltar is a viable solution despite these limitations because it uses fewer system resources and is less intrusive than current solutions. Because every attack consumes power, Gibraltar's integration of battery behavior monitoring, along with conventional IDS and anti-virus precepts, provides an additional layer of defense that is currently needed for this computing and network sector. Section III outlines the methodology for the Gibraltar DIPS, Section IV provides the complementary implementation and testing of the components that support this methodology and Section V discusses the future work and lessons learned from these test results before concluding.

## II. Methodology

**a. Determining battery activity** – Smart batteries have three operating states: Busy, Idle, and Suspend. Each of these states has a respective power usage. During the Busy state, there may be many operations running at once, so the current draw from the battery is much higher. During the Idle state, the current draw from the battery is the minimum at which a user can interact with the device [2]. During the Suspend state, data is stored in memory for future use, but no network user can interact with the PDA until a button is pressed. The Battery Signal Amplifier & Filter Engine (B-SAFE) method of measuring the battery’s current runs in either the Idle or Busy states, the Idle state being preferable to isolate any particular power signature.

In order to detect the presence of a variety of attacks, we implement the capabilities of the B-SAFE design. This testing platform allows us to clearly observe the current draw from the battery during different attacks. Using external devices, such as the Agilent 54622D Oscilloscope and a Bipolar DC Power Supply, the B-SAFE outputs a signal that is easily manipulated using AutoSignal® conditioning software. B-SAFE has two components: the Battery Harness and the Instrumentation Amplifier Circuit. The oscilloscope functions as the capture device of the amplified voltage during an attack, and AutoSignal® allows manipulation of the signal in the frequency domain with Periodogram calculations to accentuate several key frequencies that collectively constitute an attack’s signature. Figure 1 provides a diagram of the basic B-SAFE configuration.

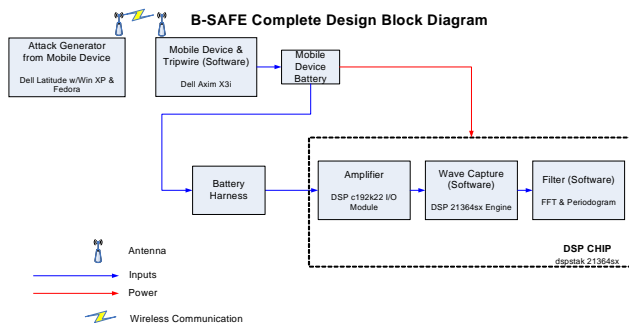


Figure 1. B-SAFE Design Flow and Major Components

In order to conduct testing, the battery harness built extends the connection between the battery and the PDA, allowing the PDA to function with the battery separated. The harness also allows us to intercept the current flowing through the PDA from the negative voltage terminal of the battery. By placing a resistor in series with the current, we create a voltage drop

(proportional to the current) that can be amplified using an instrumentation amplifier IC.

A trigger is necessary to set a threshold at which the oscilloscope captures the amplified voltage. Our trigger works in “sequence”, meaning the oscilloscope does not mistake a random spike in voltage as the actual attack. Instead, the capture initiates only after a series of repetitive voltage spikes, indicative of a repeating ping or attack on the device. As outlined in Section IV, we have successfully used this trigger to capture the voltage waveform on a variety of denial of service (DoS) attacks, such as TCP, UDP and ICMP ping floods, and non-denial of service (nonDoS) attacks, such as LSASS, and various other buffer overflow exploits.

The captured voltage waveform is then passed on to AutoSignal® for manipulation. The Fast Fourier Transform (FFT) and peak manipulation algorithms analyze the signal in the frequency domain to establish the predominant peaks into specific frequencies from the voltage pattern originally captured in the time domain. After repeating these attacks, a standard magnitude and frequency range for peaks is set and Periodogram algorithms programmed for the PDA are applied to derive and uniquely identify signatures in the frequency domain [3]. Because there is very little deviation in the magnitude and frequency of the dominant peaks using this method, a battery powered mobile device can provide a novel and mathematically powerful means to accurately detect and identify attacks using direct current. This technique is not readily available to AC powered workstations because their power supply is conditioned and constant. In effect, Gibraltar’s DIPS enables mobile devices to provide innovative and meaningful sensor-like security to alert and protect the host (as well as the network it is associated with) by detecting threshold violations in power consumption as well as detecting and identifying a variety of attacks.

**b. Attack Signatures** – Gibraltar relies on being able to recognize attack signatures on mobile devices and on the uniqueness of attack signatures. To date, the majority of attacks launched on mobile devices with B-SAFE integrated have produced unique frequency signatures. Our research has shown that even after making small changes to the code of attacks (to represent how *script kiddies* might manipulate and compile attacks downloaded from the Internet), the attack type will still be recognized.

Presently, there are a number of known mobile attacks that range from Denial-of-Service attacks to nonDoS buffer overflow exploits. In lieu of developing the

signatures for *all* known mobile attacks, we took a set of common mobile attacks [4] and determined their battery activity signatures. These attack frequency signatures populate a database in the device that can be used to compare captured signatures of suspected attacks. If no signature match is made, an extension of this process is to send the captured xy pairs of time versus voltage data from this event to a security server that contains signature patterns for a much larger set of attacks and anomalies.

**c. Result Processing** – Once an anomaly is detected based off signature analysis, additional forensics are used to enhance the view of the system status near the moment of attack or anomalous activity. Coupling the signature analysis with the forensics view of the system establishes a compelling intrusion detection process is established on the mobile device. This analysis determines if the current battery activity signifies an attack or anomalous behavior. The comparison will be logged for further analysis and then other common intrusion detection methods can be launched that will complement this type of processing scan, such as active processIDs, IP-header information of incoming traffic and changes to the registry.

**d. Communication** – The data collected by the scanning tools on the host devices provide a network administrator additional capabilities in identifying anomalous activity on mobile devices. A means of communicating this information to a network administrator seamlessly is critical for this enterprise solution. Thus, scanned data is sent to a Collection Analysis Correlation (CAC) server for further analysis, which, in turn, can send requests for these forensic-like reports proactively to other distributed mobile devices.

## IV. Design, Implementation, and Testing Results

Our test results focus on the attacks against PDAs (Dell Axims) running PocketPC 2003 and Windows Mobile 5.0 due to this operating system's increasing market share in a variety of mobile products. We use the latest versions of Visual Studio .NET 2003 along with the current version of the .NET Compact Framework (v2.0) and .NET Compact Framework SDK. We then integrate the power related structures in .NET, API member function calls, filtering algorithms as well as some of our own coding in C#. Since the utility of the .NET Framework is to make seamless integration between windows platforms [5], software can be ported over to multiple mobile architectures that support the Microsoft .NET compact Framework, such as Windows CE-based Smart Phones, PDAs or similar devices. Using this toolkit, we built Gibraltar by

creating a Signal Analysis Machine (SAM), a Digital Signal Translation and Algorithm Benchmark (DSTAB) and a Distributed Intrusion Protection System (DIPS) module. For testing purposes, we also developed *TadDah*, a functional trojan/rootkit for Windows Mobile devices.

### a. SAM

**a.1. SAM Operation** – One goal for Gibraltar is to be able to recognize attacks by their power signature after they are converted into the frequency domain. To do this the power signal must be filtered in order to compare and recognize attacks. There will always be some power being lost as long as the device is running. This underlying signal can be filtered out as well as noise in the signal after using a FFT algorithm by filtering out the lower common harmonics (below 250Hz) as well as the higher end white noise (above 2KHz.) Next the dominant peaks are derived from Periodogram calculations. Finally analysis must be done on the variance on the top five dominant xy pairs from the Periodogram. Figures 2.0, 2.1 and 2.2 provide the graphical images of how the power signature is converted into the frequency domain and then how unique xy pairs of the dominant frequencies are extracted.

**a.2. SAM Design** – SAM can identify known attacks based on their power usage signature. Knowing that an attack is in progress is helpful to the user, but knowing the specific attack being launched against a device or group of devices within a domain is even more helpful to administrators responsible for protecting the entire network. The knowledge that an attack does not fit a set of known attacks can also benefit to network administrators. The knowledge that an attack does not fit a set of known attacks can also be beneficial to network administrators. For example, when an attack is occurring, it is important to provide as much device information as possible to the network administrator in order to record its profile for (future) security and forensics purposes.

**a.3. SAM Testing Results** – SAM was tested using several known attacks. During testing, we ran attacks repeatedly to determine if their signature fell within an accepted range of variance each time in order to ascertain if they could be uniquely recognized consistently. The tests were conducted using a variety of common user configurations and settings, e.g., screen settings and common programs running. To help reduce the chance of false positives, we conducted a series of tests in each set of conditions to help ensure that normal activities, such as opening a program, were not incorrectly categorized as an attack. Although we

can not account for all permutations, addressing such instances gives us a greater degree of certainty when classifying anomalies.

Although an advantage of SAM is its ability to detect and identify different types of DoS attacks, it could conversely be subdued by an attacker using known signatures of a valid activity to mask an attack. For example, if an attacker could craft an attack to look like normal activity, they might avoid detection. However, such a feat is not trivial because the exploit would need to match both traffic and timing characteristics as well as the underlying attack activity. Another minor disadvantage is that SAM takes some system resources to filter and compare signatures. This disadvantage is mitigated however by running all of the utilities of SAM only after a threshold violation has occurred.




Figure 2.0 Example Attack Signal in Time Domain

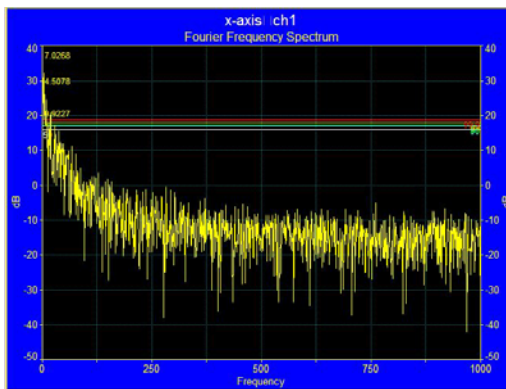


Figure 2.1 Example FFT on Attack Signal

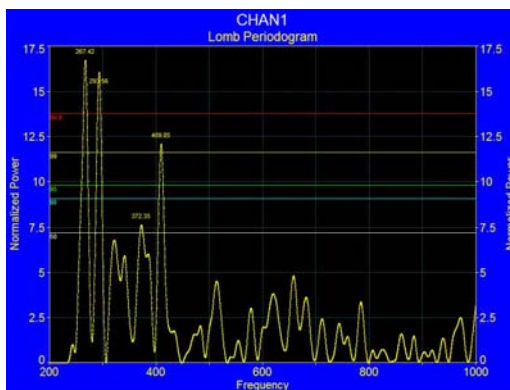


Figure 2.2 Example Periodogram from FFT of Attack

## b. DSTAB

**b.1 DSTAB Operation** – Two major paradigms in audit based intrusion detection systems are anomaly detection and misuse detection. Anomaly detection attempts to find activities that deviate from the norm. Misuse detection attempts to match activities with known malicious behavior [6]. The Digital Signal Translation and Algorithm Benchmark (DSTAB) aims to utilize the individual advantages of each of these methodologies while mitigating the potential hazards, excessive resource demands.

Limited energy within a portable device prohibits continuously running of the DSTAB suite of scan tools. Thus, execution of DSTAB forensics tools is predicated on the power usage tripwire of SAM that, upon threshold violation or user/administrator direction. DSTAB conducts the following: (SAM), process identification scan, port scan, and registry scan. In this way, DSTAB requires minimal power consumption to detect anomalies that could be malicious. Upon threshold violation or user command, DSTAB initiates the SAM component to compute data signals processing. Once DSTAB detects anomalies in device usage, it initiates one or more scans in the suite to detect known misuses. DSTAB then determines the result locally on the host or, if necessary, transmits an SML string of the information it has acquired relating to this incident so a more complete picture can be provided by the CAC server. These four tools are not completely novel in their application; however, working together they provide a unique and powerful forensics tool on mobile devices. Individually, the four scans provide a limited view of the device state. However, individual host reports in the aggregate give even more insight regarding the state of devices on the network and the network as a whole.

**b.2. DSTAB Design** – DSTAB can provide communication with the global system and enforce user defined or global network policies. It is also capable of operating as a standalone intrusion protection system. However the true power of DSTAB is attained as part of a globally distributed system. DSTAB operates as a host-centric intrusion protection system capable of acting as a seamless node within a network-centric defense system. As it feeds back reports that may provide earlier warning of an attack on network misuse.

**b.3. DSTAB Testing Results** – A reasonable estimate of power consumption can be derived from a device in the Idle state. Additionally, systematic calculations can stabilize a threshold over time. However, the device may be vulnerable if an attack takes place during startup, as it skews the threshold value. This limitation

is mitigated by a startup scan of the device to search and resolve any abnormalities or known malicious activities. Figure 3 shows several DSTAB interfaces that can be user generated or automated to provide activity data for a forensics log for proactive reporting and aggregate analysis.



Figure 3. DSTAB Event Data Pull Interfaces

### c. DIPS

**c.1. DIPS Operation** – The Distributed Intrusion Protection System module provides the means to communicate anomalous and/or possibly malicious activity to a Collection Analysis Correlation server. Scan data (to include the activity signatures from SAM and device data from DSTAB) moves to the centralized CAC server via internet protocols. The server can quickly collate and correlate reports from multiple mobile devices. When a malicious activity is discovered, the administrator can disseminate security measures back to the devices. The dissemination is a reactive and proactive response, which provides a faster and more comprehensive means of reacting to exploit vulnerabilities and other malicious activity. This reactive and proactive response in an enterprise application provides an enhancement to the security on mobile devices and, in effect, the enterprise network.

**c.2. DIPS Design** – DIPS transmits the log file via UDP or TCP wireless connection from the device to the CAC server. The scanning data is compiled into an XML message string and collected and compiled for analysis by the CAC server. The data is parsed using common parsing techniques and correlated with other

data from other devices. Based on the results of the analysis, any protective actions that the network administrator would take are sent back to the client device.

**c.3. DIPS Testing Results** – Although less prone in a TCP connection, the UDP transfer method of a log file experiences corruption if the wireless connection is not ideal. This corruption is more likely in the transmission of large files (e.g., greater than 1MB). Consideration must be placed on the size of the log file to limit corruption, as well as a method for requesting to retransmit a log file that was corrupted. These circumstances can be favorable in that it enforces a log policy on size constraints; thus, limiting the device's resource utilization for the log transfer and conceivably more reliable transmission because of smaller file size. In addition, UDP is more suitable for the asynchronous nature of wireless devices. It is more appropriate to handle infrequent corruption, than rely on a TCP connection that may not always be possible.

### d. TadDah

**d.1. TadDah Operation** – TadDah is a small, malicious, client-server application that is capable of running undetected by typical users and automatically loads at system start. The program provides a remote user with the functionality to conduct malicious activity that is comparable or superior to known mobile device trojans and rootkits. The sever can run standalone or it can be integrated into seemingly useful software. TadDah can run on any device that supports the .NET compact framework.

**d.2. TadDah Design** – TadDah has the functionality to push and pull files between devices, kill processes, capture device system info, freeze the device and force a reset, capture device screens and display messages to the user. TadDah operates on port 8000 and the actual process name is PPCServer.exe TadDah integrates characteristics that are common in PC rootkits that have not yet been developed for mobile device variants.

**d.3. TadDah Testing Results** – Gibraltar can correlate on three different sources of information to discover TadDah. Gibraltar detects TadDah running via process analysis that is unavailable through the local device process viewer. Additionally, TadDah is present in a registry field. Gibraltar also detects port 8000 as being open. Gibraltar can also provide the IP address if the attacker is connected to the device and log the information. The structure of Gibraltar allows for modularized policy additions on the local host via the centralized server. Tests with TadDah strongly suggest

Gibraltar is capable of detecting and correcting malicious software in the realm of rootkits and trojans.

## V. Future Work

In its current form, Gibraltar is mainly limited by today's commercially available Smart Battery. The Smart Battery provides the means in which we can extrapolate battery activity from the device. These batteries have an embedded chip that operates at 1 Hz – extremely slow when compared to the equipment we used to determine attack frequency patterns from power signatures. Therefore, to increase the performance of Gibraltar's scanning capability, a more robust embedded chip to measure the battery activity is necessary. Future work on Gibraltar would benefit from the development of a Digital Signal Processing (DSP) chip capable of faster processor speeds, which would allow the accuracy and precision of B-SAFE. The DSP chip would also be programmable so that forensics operations in software could be moved to the hardware level [7]. This would dramatically increase the speed and performance of Gibraltar's capabilities – in addition to taking less of a toll on the battery. The B-SAFE design and operation are presently being used to build such a DSP.

Another improvement to the DSTAB functionality would be incoming packet analysis. In addition to simple packet analysis of incoming/outgoing traffic, DSTAB could do a packet dump of all detected traffic via promiscuous mode. Consequently, DSTAB could log all detectable wireless access points (WAP), similar to *WIFIFoFum* or *MiniStumbler*. Correlation of WAP and known IP ranges could prove useful to administrators where an attacker may be using a rouge access point to launch attacks or to determine which WAP the device is connected to. Using these technologies to further compliment the forensics tools of DSTAB would enable Gibraltar to process and identify anomalous activity with greater accuracy and even pinpoint the geographic origins of the attack.

## VI. Conclusion

The effectiveness of Gibraltar's functionality outlined in this paper demonstrates a security approach for mobile devices that efficiently utilizes resources to provide an intrusion detection and protection system. Compared to common security methods, namely on the PC platform, Gibraltar takes a novel approach to intrusion detection, capitalizing on the Smart Battery's features and complementing its reports with a variety of data gleaned from event activities and the system registry. Currently a performance hindrance on mobile devices, the battery can also provide a security strength

that provides both efficient and effective detection system for a variety of attacks and anomalies. Gibraltar's approach to detection and identification enables mobile devices to serve as viable sensors for virtually any network defense strategy. With such a system in place, the most vulnerable network device can become an invaluable earlier warning and response platform.

## References

- [1] McAfee AVERT Labs Internet WWW page, at URL: < <http://phx.corporate-ir.net/phoenix.zhtml?c=104920&p=irol-newsArticle&ID=796926&highlight>> (last accessed 03/06/2006).
- [2] Grant A. Jacoby, Randy Marchany and Nathaniel J. Davis IV, "Battery-Based Intrusion Detection: a First Line of Defense," *Information Assurance Workshop 2004*, June 2004.
- [3] Grant A. Jacoby and Nathaniel J. Davis IV, "Battery-Based Intrusion Detection," *GlobeComm 2004*, December 2004
- [4] "The Twenty Most Critical Internet Security Vulnerabilities", SANS Institute, Internet WWW page, at URL: < (<http://www.sans.org/top20/>).> (last accessed 01/19/2006).
- [5] Microsoft Corporation, Advanced Power Management The Next Generation. Version 1.0. Internet WWW page, at URL:< [http://www.microsoft.com/whdc/hwdev/archive/busbios/amp\\_12.msp](http://www.microsoft.com/whdc/hwdev/archive/busbios/amp_12.msp) > (last accessed 12/24/2005).
- [6] J. Cannady, J. R. Harrell, J.R. "A Comparative Analysis of Current Intrusion Detection Technologies". *Proceedings of Technology in Information Security Conference (TISC)*, pp. 212-218, 1996.
- [7] Koopman, P., Embedded Security, *IEEE Computer*, pp. 95-97, Jul. 2004.