

# Notary

## A Device for Secure Transaction Approval

ANISH ATHALYE, ADAM BELAY, M. FRANS KAASHOEK, ROBERT MORRIS,  
AND NICKOLAI ZELDOVICH



Anish Athalye is a PhD student in the PDOS group at MIT, working on systems, security, and formal verification. [aathalye@mit.edu](mailto:aathalye@mit.edu)



Adam Belay is an Assistant Professor of Computer Science at MIT's Electrical Engineering and Computer Science (EECS) department, and a member of the Computer Science and Artificial Intelligence Lab. He received a PhD from Stanford for his work on high performance networking. Recent projects include Shenango, an operating system that improves datacenter efficiency, and Shinjuku, a system that uses fine-grained preemption to reduce tail latency. His current research focuses on the intersection of hardware and software, with an emphasis on improving security and performance. [abelay@mit.edu](mailto:abelay@mit.edu)



Frans Kaashoek is the Charles Piper Professor in MIT's EECS department and a member of CSAIL, where he co-leads the Parallel and Distributed Operating Systems Group (<http://www.pdos.csail.mit.edu/>). Frans is a member of the National Academy of Engineering and the American Academy of Arts and Sciences, and the recipient of the ACM SIGOPS Mark Weiser award and the 2010 ACM Prize in Computing. He was a cofounder of Sightpath, Inc. and Mazu Networks, Inc. His current research focuses on verification of system software. [kaashoek@mit.edu](mailto:kaashoek@mit.edu)

**N**otary is a new design for a hardware wallet, a type of security token that is used to protect sensitive transactional operations like cryptocurrency transfers. Notary aims to be more secure than past hardware wallets by eliminating classes of bugs by design and by formally proving the correctness of the key operation used in its implementation. We built a physical prototype of Notary and showed that it achieves functionality similar to existing hardware wallets while avoiding many bugs that affect them.

Hardware wallets, USB keys with a display, buttons, and the ability to run custom code, aim to provide a secure platform for approving transactions such as bank transfers and cryptocurrency transactions. By moving security-critical approval decisions to the device, hardware wallets remove the need to trust relatively complex and bug-prone computers to achieve overall application security. Hardware wallets run multiple applications, which need to be isolated from each other. Existing wallets do this using a traditional operating system design that relies on hardware protection mechanisms like CPU privilege levels and memory protection, but, unfortunately, existing wallets suffer from bugs similar to those that plague traditional computer operating systems.

Notary is a new hardware wallet that aims to avoid many of these bugs by design. Notary achieves strong isolation using *reset-based switching*, along with the use of a physically separate system-on-a-chip for running untrusted code. Notary has a machine-checked proof of the hardware's register-transfer level (RTL) design and software, showing that reset-based switching leaks no state between applications. We built a hardware/software prototype of Notary, along with a number of apps that run on the device, and demonstrated that Notary's design avoids many bugs that affect past hardware wallets.

### The Hardware Wallet Paradigm

Users routinely rely on their computers or smartphones to perform and approve security-critical operations. These operations include financial operations, such as bank transfers and cryptocurrency transactions, and non-financial operations, such as system administration tasks like deleting backups or modifying DNS records. The security of these operations relies on the security of the application as well as the underlying platform. Unfortunately, modern computers are inadequate for this purpose because they have complicated software stacks that are full of bugs; even smartphones, often thought to be more secure than PCs, have fallen victim to jailbreaks and malware. On these platforms, buggy or malicious applications might tamper with security-critical operations. Is it possible to achieve security for sensitive transactional operations even when the PC and smartphone are compromised?

Recently, we have seen an increase in the adoption of two-factor authentication (2FA) devices such as Universal 2nd Factor (U2F) tokens, devices that usually come in the shape of a small USB stick and augment the PC to provide additional security for logins. However, these 2FA devices are a bit of a red herring when we are worried about the security of the platform itself, because 2FA devices authenticate the *login* process but not the rest of the interaction with the application. This helps defend against a certain class of attacks, such

## Notary: A Device for Secure Transaction Approval



Robert Morris is a Professor in MIT's EECS department and a member of the Computer Science and Artificial Intelligence Laboratory. He received a PhD from Harvard University for work on modeling and controlling networks with large numbers of competing connections. His interests include operating systems and distributed systems. [rtm@csail.mit.edu](mailto:rtm@csail.mit.edu)



Nikolai Zeldovich is a Professor of EECS at MIT and a member of the Computer Science and Artificial Intelligence Lab. He received his PhD from Stanford University in 2008. Recent projects by Prof. Zeldovich and his students and colleagues include the CryptDB encrypted database, the STACK tool for finding undefined behavior bugs in C programs, the FSCQ formally verified file system, the Algorand cryptocurrency, and the Vuvuzela private messaging system. His current research lies in building practical verified systems. [nickolai@csail.mit.edu](mailto:nickolai@csail.mit.edu)

as a stolen password: an attacker would not be able to log in to a victim account without the second factor. But it does not help when the platform is compromised: malware on a user's computer waits until the user logs in to a target service (using their U2F token), and then the malware uses the valid session to perform malicious actions.

In contrast, hardware wallets can provide security even when the user's computer is compromised. In the hardware wallet paradigm, an application is refactored to separate out security-critical approval decisions from the rest of the application. An untrusted part of the application runs on the user's PC, while a trusted security-critical *agent* runs on the hardware wallet and is used for approving transactions. The wallet has a display where it shows the user a transaction, and it has buttons to allow the user to confirm or deny the transaction. The approval is required to go through the hardware wallet, and this is generally enforced by requiring a signature with a private key that's stored only in the wallet.

Cryptocurrencies already fit this paradigm where the approval decision is cleanly separated out, and in fact, hardware wallets are already popular with users of cryptocurrencies. For example, users run Bitcoin wallet software on their PC, where they can view their balance, view past incoming and outgoing transactions, and set up transfers, but they cannot actually transfer currency. To send bitcoins, the user crafts a transaction on their PC and sends it to their hardware wallet, which parses the transaction and displays on its screen a human-readable description like "send 1.3 BTC to 1M3K...vUQ7." Only if the user presses a "confirm" button on the hardware wallet does the device sign the transaction, which enables it to be processed by the Bitcoin network.

The paradigm of authenticating transactions on a separate, secure device has gained traction among cryptocurrency users, perhaps due to the high-stakes nature of irreversible transactions. The idea has not yet caught on with more traditional client-server applications like web apps, but there has been some progress in that direction. For example, the Web Authentication API has an extension for transaction authorization, which allows for displaying a prompt string on an authenticator device and receiving confirmation from the user [1].

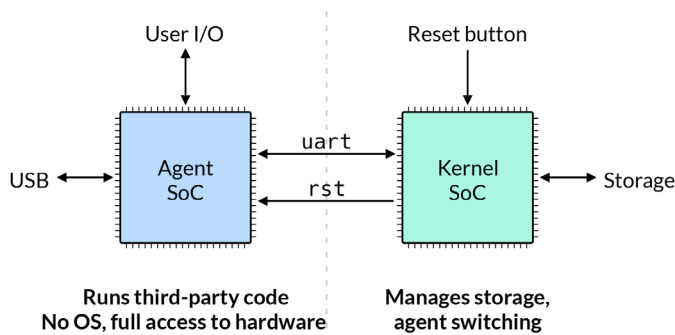
### ***Hardware Wallets Can Have Bugs Too***

With hardware wallets, the PC is removed from the trusted computing base: security depends only on the wallet, which is a big win in terms of security. These devices are much simpler than PCs, and the belief is that while the PC may have been difficult to make secure, the simplicity of wallets allows for more secure designs.

Most hardware wallets today are fixed-function, in the sense that they don't run third-party code: they have built-in support for some fixed set of agents, for example a particular set of cryptocurrencies, and users depend on the firmware vendor to add support for specific applications. This has the obvious downside in terms of usability: when new applications come out, such as a new cryptocurrency, users have to hope that the device manufacturer implements support. The developer of the cryptocurrency has no power to add the support themselves. On the other hand, high-end wallets on the market, such as the Ledger wallet [2], support downloading and running multiple third-party agent applications on the device. This is great for usability, but it adds considerable complexity, requiring that the device be capable of isolating agents from each other, because these third-party agents could be buggy or malicious.

Current devices achieve this by multiplexing the shared hardware between mutually untrusting agents with a traditional operating system using hardware protection mechanisms like CPU privilege modes and memory protection. This leads to the potential for the same kinds of bugs that exist in PC operating systems. And, indeed, existing hardware wallets have suffered from isolation bugs in memory protection configuration, system call

## Notary: A Device for Secure Transaction Approval



**Figure 1:** Notary’s design physically separates trust domains with an SoC per domain and a simple interconnect between trust domains (reset wire and UART). Untrusted programs are run one-at-a-time on the Agent SoC, which has its state cleared between running agents.

implementations, and driver code [3, 4]. There is also potential for hardware-related bugs: any shared hardware state could potentially be used to infer information about other applications (this is what is happening in attacks like Spectre, for example).

### Notary’s Approach

Notary is a hardware wallet that aims to avoid by design many of the security issues that affect past wallets. Notary doesn’t rely on hardware protection mechanisms like CPU privilege modes or memory protection, and it doesn’t have any system calls or even an operating system in the traditional sense. Instead, Notary is built around the idea of achieving isolation by using a dedicated system-on-a-chip (SoC), with its own CPU and memory, to run untrusted programs. Notary runs one program at a time on this chip, and it completely resets this chip (and all of its internal state) when switching between programs, a primitive that’s formalized and proven correct in our prototype. Running untrusted code on the dedicated SoC is orchestrated by a separate chip that never runs third-party code.

Figure 1 illustrates Notary’s design. The design is structured around physical separation. Notary consists of two security domains, each with its own separate system-on-a-chip (SOC), which includes a CPU, ROM, RAM, and peripherals such as UART. One domain runs the kernel, and one domain runs third-party agent code. The Kernel SoC manages persistent storage and switching between agents; no third-party code ever runs on the Kernel SoC. The Agent SoC, which has no mutable non-volatile storage, runs agent applications one-at-a-time directly on raw hardware (with no OS to protect the hardware). The Agent SoC has direct access to the user I/O path, the buttons and display, as well as access to USB to communicate with the outside world.

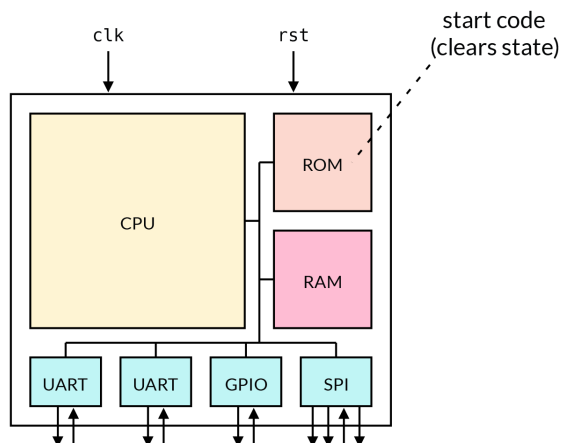
In this architecture, after the user chooses an agent to run, it is launched as follows. First, the Kernel SoC resets the Agent SoC and clears all of its internal state. Next, the Kernel SoC reads an

agent’s code, keys, and data from persistent storage and sends it over the UART; on the other side of the UART, the Agent SoC’s bootloader receives the code/data, saves it in RAM, and executes it. At this point, the agent runs directly on top of the hardware on the Agent SoC, not requiring further interaction with the Kernel SoC. The agent has access to everything it needs: its own code and data, the user I/O path, and communication to the outside world. It can do its job, such as displaying a Bitcoin transaction, receiving confirmation from the user, and sending a signed transaction out via USB. Finally, when the agent is done, it has only one way of interacting with the Kernel SoC: a “save and exit” operation, where the agent requests termination, optionally supplying a new persistent state. After this, to run a different agent on the device, the process starts over, beginning with clearing state in the Agent SoC. Notary’s separation architecture has analogs for all the operations that hardware wallets generally support: factory-resetting the device, installing/removing agents, and launching agents.

In Notary’s design, the decision to connect user I/O and USB directly to the Agent SoC is important for security. An alternative design might connect these to the Kernel SoC, but that would be undesirable because it would introduce the need to have communication between the Agent SoC and Kernel SoC during regular agent operation, adding complexity by requiring a large number of system calls beyond the single save/exit “system call” that Notary supports.

In Notary’s design, it is safe to give untrusted code raw access to the user I/O and USB peripherals because the state clearing operation covers peripherals: if a malicious or buggy agent puts the display or USB controller into a bad state, the reset and state clearing operation will fix it. Furthermore, having the display connected to the Agent SoC running potentially untrustworthy code does not introduce the possibility of confusing the user, due to Notary’s reset-based workflow. The user switches applications by restarting the entire device, which makes the kernel start a special agent, the application launcher, on the Agent SoC. The user can unambiguously select an agent to run, and after that point, the chosen agent has exclusive control over user I/O until the device is restarted.

With this architecture, Notary achieves isolation between two agents running one after another on the same chip. Running agent code directly on top of raw hardware, using reset as a mechanism to switch agents, obviates the need for a traditional operating system and hardware protection mechanisms, which can be error-prone to program. Performing state clearing, wiping out all state in the Agent SoC between running different agents, ensures that one agent’s secrets can’t leak to another. Essentially, Notary boils isolation between agents down to state clearing.



**Figure 2:** A schematic of Notary's Agent SoC. Carefully written code in boot ROM clears all internal state in the SoC after reset.

### State Clearing

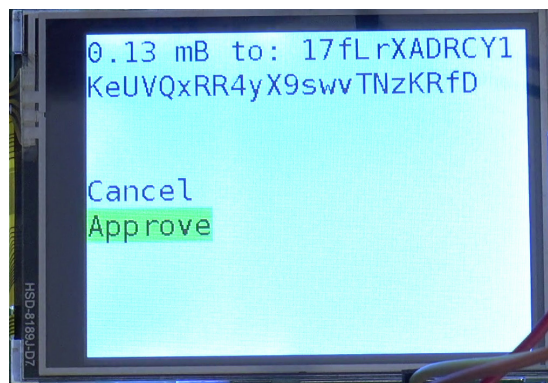
Clearing all internal state in a SoC turns out to be challenging, and simple approaches don't work.

At first, we thought that asserting the reset line of an SoC might be adequate. It turns out that ISAs don't guarantee that reset clears internal state; for example, the RISC-V ISA says that the program counter is set to an implementation-defined reset vector, and all other state is undefined [5]. In practice, many chips implement reset such that it only does the minimal work necessary to get the chip going again. For example, on our SoC, asserting the reset line did set the program counter to a well-known value, but it left much state inside the SoC untouched, including in registers, some CPU-internal caches, RAM, and peripherals.

Another approach we considered is power cycling the SoC to clear its internal state. However, research has shown that state inside these chips can persist for minutes without power [6]. Notary applies state clearing before every application switch, so a delay of several minutes to clear state would translate to a delay of several minutes when launching any application, making the device unusable. Furthermore, powering off the SoC for a few minutes provides no guarantees that state is actually cleared.

### Provably Correct Software-Based State Clearing

Notary's approach is to use software to clear an SoC's state. The idea is that asserting the reset line resets the program counter, so it can return control to software in boot ROM that can complete the job of clearing all state in the chip, as shown in Figure 2. The idea of having initialization code run on startup is not new, but Notary's boot code is doing something unusual: it's aiming to clear every bit of state *internal* to the SoC, which includes details



**Figure 3:** Notary prototype running a Bitcoin wallet agent

that don't even exist at the ISA level, such as microarchitectural state. Writing this boot code is a challenge; it's not immediately obvious that writing such code will even be possible. We normally think about code at the abstract machine level, consulting the ISA specification to understand its behavior, but in Notary's case, we need this code to affect internal state.

To help develop this boot code and convince ourselves that it's correct, we built a tool that analyzes an SoC's implementation at the gate level to determine whether the boot code successfully clears all internal state in all situations. The tool takes Verilog code that describes the SoC, converts it to a format compatible with SMT solvers, and then checks whether boot code running on the chip satisfies our correctness property by simulating the circuit symbolically.

Notary's boot code for its simple RISC-V-based SoC, built on the PicoRV32 [7], is formally verified to clear all SoC-internal state correctly using this tool. We are currently working on applying this technique to more complex SoCs.

### Prototype

We built a hardware/software prototype of Notary, along with a number of agents that run on the device: a Bitcoin agent and a general-purpose web app approval agent similar to Web Authentication. Figure 3 shows our prototype running the Bitcoin agent in the process of approving a transaction. In our prototype, the heavyweight reset-based approach for launching agents takes about 135 ms, fast enough for interactive use. Of this, 7 ms are spent running the formally verified state clearing code, with most of that time used clearing RAM, and the rest spent copying the agent code/data to the Agent SoC over the relatively slow UART.

## Notary: A Device for Secure Transaction Approval

### Conclusion

Notary is a case study in designing for security. Notary simplifies software (e.g., using reset-based agent switching) and wastes resources (e.g., using physical separation) in order to achieve strong isolation and defense in depth. This separation and reset-based switching eliminates by design classes of bugs that affect traditional user/kernel co-resident designs, including OS bugs, microarchitectural side-channels, and certain hardware bugs. Notary can improve the security of applications where the crucial transaction decision can be succinctly summarized and delegated to a strongly isolated agent running on Notary.

So far, cryptocurrencies have embraced hardware wallets, with significant adoption by users. In the future, we hope to see more applications be refactored to take advantage of the enhanced security that hardware wallets offer.

The full Notary paper is available at <https://pdos.csail.mit.edu/papers/notary:sosp19.pdf>.

### References

- [1] W3C, “Web Authentication: An API for Accessing Public Key Credentials,” March 2019: <https://www.w3.org/TR/webauthn>.
- [2] “Ledger Hardware Wallets”: <https://www.ledger.com>.
- [3] Riscure Team, “Hacking the Ultra-Secure Hardware Cryptowallet,” August 2018: <https://www.riscure.com/blog/hacking-ultra-secure-hardware-cryptowallet>.
- [4] C. Guillemet, “Firmware 1.4: Deep Dive into Three Vulnerabilities which Have Been Fixed,” March 2018: <https://www.ledger.com/2018/03/20/firmware-1-4-deep-dive-security-fixes>.
- [5] A. Waterman and K. Asanovic, “The RISC-V Instruction Set Manual, Volume II: Privileged Architecture,” June 2019: <https://riscv.org/specifications/privileged-isa>.
- [6] A. Rahmati, M. Salajegheh, D. E. Holcomb, J. Sorber, W. P. Burses, and K. Fu, “TARDIS: Time and Remanence Decay in SRAM to Implement Secure Protocols on Embedded Devices without Clocks,” in *Proceedings of the 21st USENIX Security Symposium*, 2012, pp. 221–236: <https://www.usenix.org/system/files/conference/usenixsecurity12/sec12-final71.pdf>.
- [7] C. Wolf, “PicoRV32—A Size-Optimized RISC-V CPU,” 2019: <https://github.com/cliffordwolf/picorv32>.