

Article

Not peer-reviewed version

Towards Secure and Reliable Runtime Patching for Cyber-Physical Systems in Industrial Networks

[Zesheng Xi](#), [Bo Zhang](#)^{*}, [Yunfan Wang](#), [Chuan He](#)

Posted Date: 16 May 2025

doi: 10.20944/preprints202505.1282.v1

Keywords: embedded security; runtime patching; industrial communication systems; cyber-physical resilience; power system reliability



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Towards Secure and Reliable Runtime Patching for Cyber-Physical Systems in Industrial Networks

Zesheng Xi ^{1,2}, Bo Zhang ^{2,*}, Yunfan Wang ^{2,3} and Chuan He ^{2,3}

¹ School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, 210014, China

² School of Cyber Science and Engineering, Southeast University, Nanjing, 211189, China

³ State Grid Laboratory of Power Cyber-Security Protection and Monitoring Technology, State Grid Smart Grid Research Institute Co., LTD, Nanjing, 210003, China

* Correspondence: zhangbo6@epri.sgcc.com.cn

Abstract: Embedded control devices within power system communication infrastructures are increasingly vulnerable to cyber threats due to escalating software complexity and pervasive network exposure. Conventional patching mechanisms—whether static or dynamic—often fail to reconcile the competing demands of real-time responsiveness and computational efficiency in resource-constrained environments typical of power grids. To tackle this issue, we propose a hardware-assisted runtime patching framework tailored for embedded systems in critical power system networks. Our method integrates binary-level vulnerability modeling, execution-trace-driven fault localization, and lightweight patch synthesis, enabling dynamic, in-place code redirection without disrupting ongoing operations. By constructing a system-level instruction flow model, the framework leverages on-chip debug registers to deploy patches at runtime, ensuring minimal operational impact. Experimental evaluations within a simulated substation communication architecture reveal that the proposed approach reduces patch latency by 92% over static techniques, while incurring less than 3% CPU overhead. This study offers a scalable and real-time model-driven defense strategy that enhances the cyber-physical resilience of embedded systems in modern power systems, contributing new insights into the intersection of runtime security and grid infrastructure reliability.

Keywords: embedded security; runtime patching; industrial communication systems; cyber-physical resilience; power system reliability

1. Introduction

Embedded systems form the operational backbone of modern power infrastructures, supporting critical functions in smart substations, edge controllers, and renewable integration systems [2]. However, their increasing connectivity within industrial power networks has significantly expanded the system attack surface. Recent ICS-CERT reports reveal that over 68% of industrial control systems—including those deployed in power automation environments—contain exploitable vulnerabilities [21]. These security flaws enable sophisticated cyberattacks ranging from data exfiltration to direct manipulation of grid operations, potentially causing financial losses exceeding 10 million per incident in mission-critical scenarios [22,23]. This reality underscores the urgent need for lightweight, real-time, and robust runtime protection strategies tailored for embedded devices in next-generation power systems.

Existing remediation strategies adopt either static or dynamic approaches. Static patching[4] requires complete system downtime for firmware replacement, while dynamic methods[5] introduce substantial runtime overhead (typically 15-20% CPU utilization). Both prove impractical for embedded devices constrained by:

1. **High risk of service interruption:** The static patching method requires complete downtime to complete the vulnerability repair, and for mission-critical scenarios (e.g., medical monitoring,

industrial production [15], etc.) running in embedded devices, service interruptions may lead to serious consequences, which is not conducive to system continuity assurance.

2. **High resource overhead:** dynamic patching approaches are usually designed for high-performance devices, and embedded devices are difficult to support the operation of complex patches due to resource constraints [3] (e.g., limited memory, storage, and computational power). In addition, dynamic patching may introduce additional computation and storage overheads, further affecting device performance.

To address the above limitations, this paper proposes a hot-patching-based vulnerability remediation method for embedded devices, which is based on the built-in functionality of the hardware to perform hot-patching of embedded systems, retaining strict real-time constraints while minimizing additional resource usage [14]. The method consists of four steps: vulnerability identification; vulnerability localization; patch preparation; and patch switching. The key idea is to replace instructions at runtime using an on-board debugging unit, which is readily available on commercial off-the-shelf processors. The methodology enhances the detection-based hotfixing approach without the need for binary detection and other software layers.

The main contribution of this paper:

1. proposes a customized remediation framework to meet the heterogeneous needs of embedded systems, designs a four-step remediation process (vulnerability identification, vulnerability localization, patch preparation, and patch switching) tailored to the characteristics of embedded devices, and enhances the detection-based hotfixing approach to avoid relying on the binary detection and other software layers, and improves the accuracy and adaptability of the patches.
2. In order to solve the service interruption problem caused by static patching, the framework proposes a hot-patching-based vulnerability remediation method for embedded devices, which utilizes the hardware built-in function to execute hot patches, avoiding device downtime and system restart, and guaranteeing system continuity in mission-critical scenarios.
3. In order to overcome the high overhead of dynamic patching, the framework uses on-board debugging units to dynamically replace instructions at runtime, which greatly reduces resource consumption during the repair process while maintaining the real-time nature of the system.
4. The experimental results show that the hot-patch-based vulnerability repair method proposed in this paper performs well in embedded device vulnerability repair, effectively reduces system downtime and service interruption, and provides strong support for system stability and availability. Through in-depth analysis and optimization, this paper provides an innovative and efficient solution for security maintenance of embedded devices, which has important theoretical significance and practical value.

2. Related Works

With the continuous development of Internet technology, the traditional patching method that requires constant recompilation and deployment no longer meets the rapid iteration needs of enterprises. In order to solve the problems and security vulnerabilities that may occur during the actual use of software products, and to improve the reliability and security of the system through real-time updating and repairing, as well as to reduce the maintenance cost and improve the development efficiency, the academia and the industry have turned their attention to hot patching. Jiang Wen[6] et al. proposed a hot-patching technique for software maintenance in VXWORKS operating system, using ClearCase and ICP-CI to produce hot-patches, introducing the advantages of real-time fixes and improved system reliability. Christian Niesler[1] et al. proposed a framework called HERA to hot-patch embedded real-time systems without violating real-time constraints. HERA utilizes the hardware features of commercial Cortex-M microcontrollers to perform hot patching without any hardware modifications, preserving hard real-time constraints while keeping additional resource usage to a minimum. Zhou M[7] et al. proposed a reliable real-time patching framework, called RLPatch, which allows for closed-source patching of Programmable Logic Controllers (PLCs) without the need

to access the source code. Yi He[8] et al. proposed RapidPatch, which consists of a patch generation module for generating device-specific patches from eBPF patches and configurations, and a patch validation module for checking for malicious behaviors in the patches. Mingquan Wang[9] discusses the remote hot deployment technology based on SpringBoot. The principles of Devtools and Arthas, two Java implementations of hot deployment, are described in detail, including hot deployment methods based on ClassLoader and Java Agent mechanisms. Holmbacka[10] et al. introduce a lightweight runtime update framework that replaces C-based software in embedded systems without interrupting the system. parts of the C-based software in embedded systems without interrupting the system. This framework is based on FreeRTOS and adds mechanisms for dynamically linking and relinking FreeRTOS tasks to the system. Ramaswamy[10] et al. present a new hotfix framework, Katana, which provides an effective hotfix solution for ELF binaries, which is important for improving software security and reliability. Naeem et al.[25] proposed a hierarchical distributed proxy model based on deep learning, which utilizes the capabilities of deep neural networks and combines distributed proxy architecture to achieve efficient detection and risk assessment of network intrusions. Payer[12] et al. introduce a dynamic software update mechanism that a dynamic software update mechanism that dynamically applies security patches without modifying the binary application. This mechanism combines sandboxing techniques to protect system integrity and improves application availability through a dynamic update mechanism. Chen[13] et al. propose a scheme that provides real-time update capabilities through virtualization, allowing patches and upgrades to be applied at any time without the need to be in a silent state. The advantages of this approach are its wide range of applications and good portability, which makes it particularly suitable for integration into general-purpose virtualized systems. Altekari[16] et al. explored the development and application of the OPUS tool, which is capable of dynamically applying fixes to C programs, especially in safety-critical applications that require frequent updates.

While the above hot-patch based patching approach addresses the hardware architecture constraints of IoT devices and satisfies the embedded device real-time problem, it is difficult to generate patches based on different types of devices due to the fact that vendors have a variety of types of fragmented devices, which hampers the installation of patches on devices that require high real-time [26] and high availability. In view of the above, this paper performs hot patching of embedded systems based on the built-in functionality of the hardware, which preserves the strict real-time constraints while addressing the compatibility issues.

3. Methodology

Hot-patch based vulnerability patching method gives a non-disruptive deployment update method for embedded devices based on ARM processor with built-in debugging function, which mainly contains four steps: (1) vulnerability identification; (2) vulnerability localization; (3) patch preparation; and (4) patch switching.

3.1. Vulnerability Identification

To enable structured vulnerability identification, the latest update report released by the PLC vendor is first analyzed to extract key vulnerability information. Let the report set be denoted as $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$. Named Entity Recognition (NER) is applied via a model \mathcal{F}_{NER} to extract entities containing vulnerability ID, type, affected firmware version, target controller, logic instructions, and trigger conditions:

$$\mathcal{E} = \mathcal{F}_{\text{NER}}(\mathcal{R}) = \{(id_i, type_i, version_i, ctrl_i, logic_i, trig_i)\}_{i=1}^m$$

This structured entity set provides a solid foundation for subsequent analysis, especially for identifying Major Non-Recoverable Failure (MNRF) vulnerabilities, which may lead to memory corruption and control failures with severe safety implications.

After initial extraction, the entity set \mathcal{E} is filtered and completed. Entities that satisfy the completeness condition $\text{is_complete}(e_i)$ are retained, while incomplete ones are supplemented using expert knowledge or rule-based inference $\mathcal{F}_{\text{manual}}(\cdot)$, resulting in the completed entity set:

$$\mathcal{E}' = \{e_i \in \mathcal{E} \mid \text{is_complete}(e_i) = \text{True}\} \cup \mathcal{F}_{\text{manual}}(\mathcal{E})$$

This step ensures the semantic integrity and logical sufficiency of the vulnerability data, preventing propagation of errors in later stages.

Next, vulnerable logic instructions are extracted by mapping each vulnerability to its associated control behavior using the vendor-provided PLC programming manual \mathcal{M}_{PLC} . A mapping function \mathcal{M}_{map} is applied to obtain the instruction set:

$$\mathcal{L} = \mathcal{M}_{\text{map}}(\mathcal{E}', \mathcal{M}_{\text{PLC}})$$

These instructions typically include operations for data processing, input response, and loop control, and are critical for identifying logic-level vulnerabilities that could affect system reliability[13].

Finally, based on the instruction set \mathcal{L} , trigger conditions are identified by focusing on Arabic numerals (e.g., fixed parameters) and programming logic keywords (e.g., WHILE, IF). A trigger extraction function $\mathcal{F}_{\text{trigger}}$ is defined to produce a set of conditions expressed as $t_i = f(n_k, k_j)$:

$$\mathcal{T} = \mathcal{F}_{\text{trigger}}(\mathcal{L}) = \{t_i \mid n_k \in \mathbb{N}, k_j \in \mathcal{K}\}$$

These trigger conditions may cause unsafe behaviors such as buffer overflows, denial-of-service, or infinite loops, and are crucial for modeling how vulnerabilities are activated in real-world scenarios.

3.2. Vulnerability Localization

Vulnerability localization consists of four main steps: inserting vulnerable logic instructions, setting observation points, tracing code access paths, and recording vulnerability metadata.

First, let the control program of a Programmable Logic Controller (PLC) be denoted as \mathcal{P} , composed of a sequence of logic instructions $\{I_1, I_2, \dots, I_n\}$. A specially crafted vulnerable instruction I_v is injected such that:

$$I_v = \text{Insert}(L_v, V_{\text{tag}})$$

where L_v denotes the logical location within the program, and V_{tag} is a controller-specific label that marks this instruction as vulnerable. This insertion serves to emulate or trigger latent vulnerabilities exploitable via unauthorized control actions or data manipulation [7]. Additionally, logic flag values $f \in \{0, 1\}$ can be set to modulate control flows for conditional execution, providing a lightweight mechanism to activate I_v .

Second, a watchpoint W is placed at the address of I_v using the PLC's internal debugging mechanism:

$$W = \text{SetWatchpoint}(\text{Addr}(I_v))$$

This allows the debugger to intercept and monitor any access to I_v during execution. Depending on the platform, W may be established statically in source code or injected dynamically at runtime [24].

Third, let $\mathcal{T} = \{p_1, p_2, \dots, p_m\}$ represent the set of execution traces reaching I_v . When the PLC program executes and hits the watchpoint W , the debugger logs each access path $p_i \in \mathcal{T}$ such that:

$$p_i = \text{TracePath}(I_{\text{start}} \rightarrow I_v)$$

These paths provide essential context for identifying the control flow segments that interact with the vulnerability.

Finally, all extracted vulnerability information is documented into a remapping table \mathcal{V}_{map} , recording tuples of the form:

$$\mathcal{V}_{\text{map}} = \{(FW_i, VID_i, I_v, \text{Addr}(I_v))\}$$

where FW_i is the firmware version, VID_i is the unique vulnerability ID, I_v is the logic instruction, and $\text{Addr}(I_v)$ its entry point. This serves as a reference for future remediation or automated patching processes.

3.3. Patch Preparation

Patch preparation involves three main stages: (1) integrity verification of externally received patches, (2) parsing of patch code and metadata, and (3) static analysis of patch content.

Let a received patch be denoted as \mathcal{P}_{ext} , which includes the patch code $\mathcal{C}_{\text{patch}}$ and associated meta-information $\mathcal{M}_{\text{patch}}$. Upon reception, the system computes a cryptographic hash $H(\mathcal{P}_{\text{ext}})$ and compares it to the expected hash H_{ref} :

$$\text{Verify}(\mathcal{P}_{\text{ext}}) = \begin{cases} \text{Accepted,} & \text{if } H(\mathcal{P}_{\text{ext}}) = H_{\text{ref}} \\ \text{Rejected,} & \text{otherwise} \end{cases}$$

This ensures the integrity and authenticity of \mathcal{P}_{ext} , mitigating the risk of tampered or malicious patches.

Next, the patch is parsed to extract both $\mathcal{C}_{\text{patch}}$ and $\mathcal{M}_{\text{patch}}$, where:

$$\mathcal{M}_{\text{patch}} = \{S_{\text{patch}}, L_{\text{insert}}\}$$

with S_{patch} denoting the patch size (in bytes) and L_{insert} representing the code insertion location within the original program $\mathcal{P}_{\text{orig}}$. The parser identifies additions, deletions, or modifications as:

$$\Delta\mathcal{C} = (\mathcal{C}_{\text{add}}, \mathcal{C}_{\text{del}}, \mathcal{C}_{\text{mod}})$$

A precise determination of L_{insert} is essential to apply $\Delta\mathcal{C}$ correctly; any deviation could lead to functional failures or latent bugs [21].

Finally, a static analysis function $\text{Check}_{\text{static}}(\mathcal{C}_{\text{patch}})$ is applied using tools such as PC-Lint and Coverity Prevent [6]:

$$\text{Check}_{\text{static}}(\mathcal{C}_{\text{patch}}) \Rightarrow \text{Assessment Report}$$

This step ensures that $\mathcal{C}_{\text{patch}}$ conforms to syntax rules, complies with coding standards, avoids the use of dangerous functions, and does not introduce additional vulnerabilities. It enables a preemptive quality and security assessment before the patch is committed to $\mathcal{P}_{\text{orig}}$.

3.4. Patch Switching

Patch switching includes copying the patch code to the patch preselection area of the Random Access Memory (RAM) and adding the contents of the meta-information to the corresponding data structure; adding the patch task to the scheduling program and configuring the jump instruction; updating the program to notify the hardware of the activation of the patch through an atomic instruction; monitoring the instructions executed by the CPU, comparing the current memory location of the CPU with the insertion point of the jump instruction, and triggering an interrupt when they are consistent; jumping to the patch code through the jump instruction; and returning to and executing the patch code when it is finished. Trigger an interrupt; jump to the patch code through the jump instruction; execute the patch code and return to normal execution after completion.

3.4.1. Debugging Unit FPB Configuration

Copy the patch code to the patch preselection area of RAM and add the contents of the meta-information to the corresponding data structure[20], i.e., the debug unit FPB (Flexible Priority Blocker) configuration. The FPB unit configuration requires a C library to be executed, and this library contains the following functions:

- The `fpb_Init` function is responsible for initializing the FPB unit, ensuring that the FPB unit has been loaded correctly, and preparing for subsequent task scheduling and priority management;
- The `fpb_enable` function is used to set the FPB unit to be enabled and disabled; if not set, the device will ignore any breakpoints that have been configured and enabled;
- The `enable_single_patch` function utilizes code remapping and breakpoints to enable the creation of jump instructions for control flow redirection. The jump instructions can be provided directly by the patch or generated by calculating to get the specified offset. The function sets the breakpoint at the insertion location of the jump instruction and sets the breakpoint behavior to code remapping. Atomic switches are implemented to enable hot patching by writing in a single register;
- The `load_patch_and_dispatcher` function is responsible for patch loading and preparation, and the basic scheduler copies the patch into RAM by loading and modifying it according to the given patch location.

3.4.2. Add Patch Tasks to the Scheduler and Configure Jump Instructions

The introduction of dynamic patching technology requires extensions to handle patch tasks, including loading, execution and subsequent maintenance work. By inserting jump directives, control flow can be transferred to the patched code at runtime, but it needs to be carefully configured to ensure correct execution and avoid security issues.

3.4.3. The Update Program Notifies the Hardware to Activate the Patch Through Atomic Instructions

The update program can activate patches efficiently in real time through atomic instructions, which at the hardware level usually involves cache coherence protection and prohibiting instruction reordering to ensure data consistency and correctness during the patching process.

3.4.4. Monitor the Instructions Executed by the CPU

To realize atomic switching, hardware breakpoints can be used to enable patching. Hardware breakpoints are a special debugging feature[19], which can be set to trigger a patch switch. The patch switch is triggered when the CPU executes to the set hardware breakpoint. Since hardware breakpoints are atomic instructions, the atomicity of the patch switch can be guaranteed. The process of patch switching is a sophisticated and complex operation that involves CPU instruction monitoring and interrupt handling. In this process, the CPU plays a crucial role, continuously monitoring all instructions being executed. Whenever the CPU executes an instruction, it compares the current instruction address with a preset insertion point for jump instructions. This comparison process is extremely critical. If the CPU finds an exact match between the current instruction address and the jump instruction insertion point, it immediately triggers an interrupt to pause the current instruction flow. This is an exact breakpoint hit, meaning that a patch switch is required. At this point, the CPU stops executing the current instruction, and by writing pre-configured data to the specified hardware registers, it can realize the insertion of the jump instruction and switch to patch loading and replacement. If it does not match then the execution of the next instruction continues.

3.4.5. Jump to Patch Code by Jump Instruction

As illustrated in Figure 1, once a breakpoint is triggered, the CPU immediately replaces the current instruction at address A_{curr} with a predefined patch table entry E_{patch} . Each E_{patch} can be represented as:

$$E_{patch} = (I_{jmp}, A_{target})$$

where I_{jmp} is the jump instruction, and A_{target} is the jump destination address. The CPU fetches and executes I_{jmp} from E_{patch} , thereby redirecting control flow to A_{target} and realizing seamless instruction flow redirection.

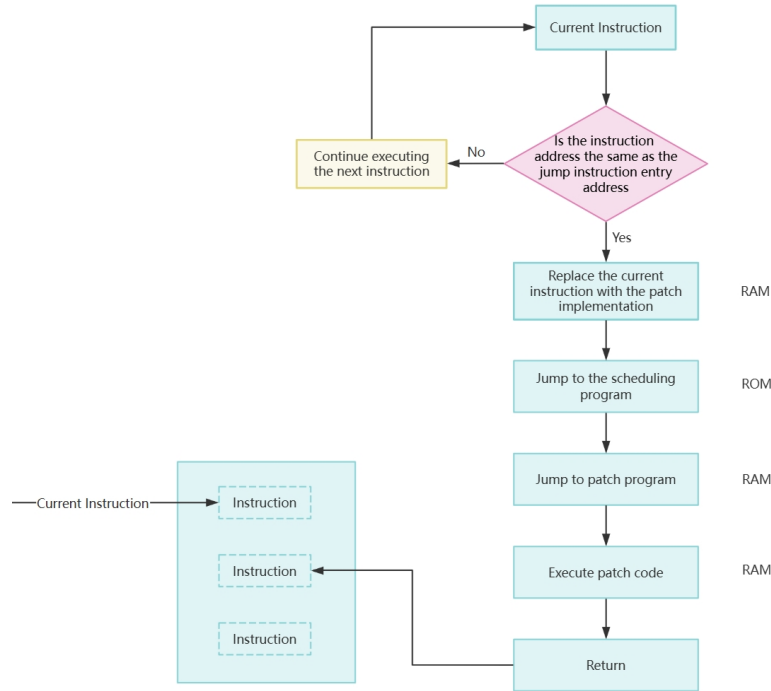


Figure 1. Flowchart of inserting jump instruction through debugging unit.

Patch injection into the original code C_{orig} is achieved by configuring two key positions: - The ****jump instruction insertion point**** $A_{\text{insert}}^{\text{jmp}} \in C_{\text{orig}}$, where I_{jmp} is inserted. - The ****branch return point**** A_{branch} , which denotes the resume point in C_{orig} after patch execution.

Thus, a redirection mechanism is constructed such that:

$$C_{\text{orig}}(A_{\text{insert}}^{\text{jmp}}) \rightarrow C_{\text{patch}} \rightarrow C_{\text{orig}}(A_{\text{branch}})$$

However, the platform in use imposes a constraint: direct jumping from ROM to RAM is not supported, i.e.,

$$\text{JMP}_{\text{ROM} \rightarrow \text{RAM}} \notin \text{ISA}_{\text{platform}}$$

To address this, an intermediate ROM region called the **jump_section** is introduced. This section acts as a ****patch dispatcher****:

$$\text{jump_section} : A_{\text{rom}} \rightarrow A_{\text{ram}}$$

It receives the control from $A_{\text{insert}}^{\text{jmp}}$, and forwards it to the patch code residing in RAM. Let D_{jmp} denote this dispatcher logic, then the complete control flow becomes:

$$C_{\text{orig}}(A_{\text{insert}}^{\text{jmp}}) \rightarrow D_{\text{jmp}}(A_{\text{rom}}) \rightarrow C_{\text{patch}}(A_{\text{ram}}) \rightarrow C_{\text{orig}}(A_{\text{branch}})$$

This mechanism effectively circumvents the hardware limitation and enables flexible patch redirection from ROM to RAM, preserving execution integrity while enabling runtime modification.

3.4.6. Execute Patch Code

When the CPU finishes executing the patch code in RAM, it returns to the original instruction stream and continues to execute the next instruction. This switching process is transparent to the user, and they hardly feel any interruption or delay. This is the essence of patch switching, which enables

fast replacement and execution of specific instructions without affecting the overall operation of the system.

4. Experiments

4.1. Experimental Environment

Table 1 shows a detailed list of experimental environment configurations.

Table 1. Experimental configuration.

Categories	Name	Description
Hardware	Multiprocessor	Intel Core i7-10870H
	GPUs	NVIDIA GeForce RTX 3060
	Random access memory (RAM)	16GB DDR4
Software	Development and operating environments	Ubuntu 20.04 LTS
	Integrated development environment (IDE)	Eclipse IDE 2021-09
	programming language	C, C++
	Debugging Tools	GDB 10.1
	version control	Git 2.25.1
	network tool	Wireshark 3.4.8
	Traffic Tools	Scapy 2.4.5
	virtualized environment	VMware Workstation Pro

4.2. Data Sets

In this study, the vulnerability fixing dataset shown in Table 2 below was used, with each vulnerability corresponding to a specific component and description, as well as a corresponding filter.

Table 2. Data sets.

Vulnerability ID	subassemblies	descriptive	(machine) filte
2020-17441	PicoTCP	Validate IPv6 payload length field against actual size for function pico_ipv6_process_in	Filter Patch
2020-17442	PicoTCP	Validate hop-by-hop IPv6 extension header length field for function pico_ipv6_process_hopbyhop	Filter Patch
2020-17443	PicoTCP	Restrict that echo->transport_len is no less than 8 in pico_icmp6_send_echoreply	Filter Patch
2020-17444	PicoTCP	Check possible overflow of header extension length field for function pico_ipv6_check_headers_sequence	Filter Patch
2020-17445	PicoTCP	Validate optlen using a loop prior to function pico_ipv6_process_destopt	Filter Patch

4.3. Experiment and Result Analysis

4.3.1. Vulnerability Fixing Time Comparison Experiment and Result Analysis

In the vulnerability remediation process for embedded devices, remediation time is a key metric. In order to evaluate the advantages of hot-patching methods in terms of fixing time, we conducted comparative experiments covering three cases: no patch added, traditional patching method and hot-patching method. The experimental results are shown in Figure 2.

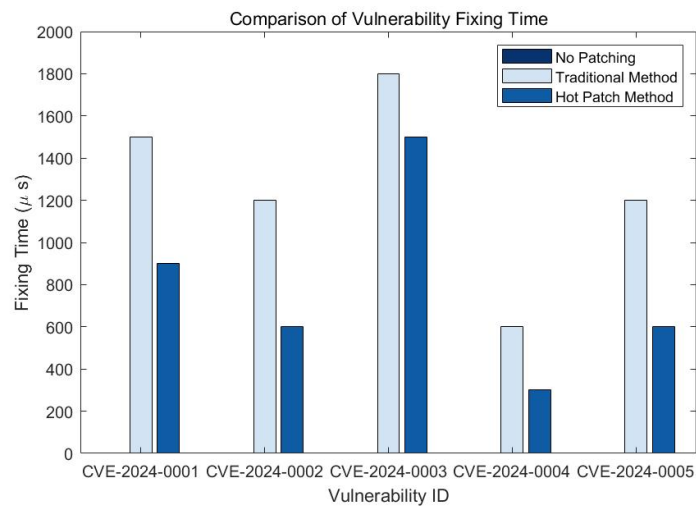


Figure 2. Comparison of vulnerability fixing time.

As seen in Figure 2, the hot-patching approach has significantly shorter remediation times than the traditional update approach for all tested vulnerability IDs. This significant difference indicates that the hot-patching approach has a clear advantage in reducing system downtime. Shorter remediation time is critical for improving system availability and minimizing service interruptions. Traditional update methods require long system downtime, which challenges the real-time and reliability of embedded devices, while hot-patching methods effectively solve this problem, allowing the system to remain running efficiently during the vulnerability remediation process.

4.3.2. Vulnerability Repair Success Rate Comparison Experiment and Result Analysis

In order to further evaluate the effectiveness of different restoration methods, we compared the performance of different methods in terms of restoration success rate. The experimental results are shown in Figure 3.

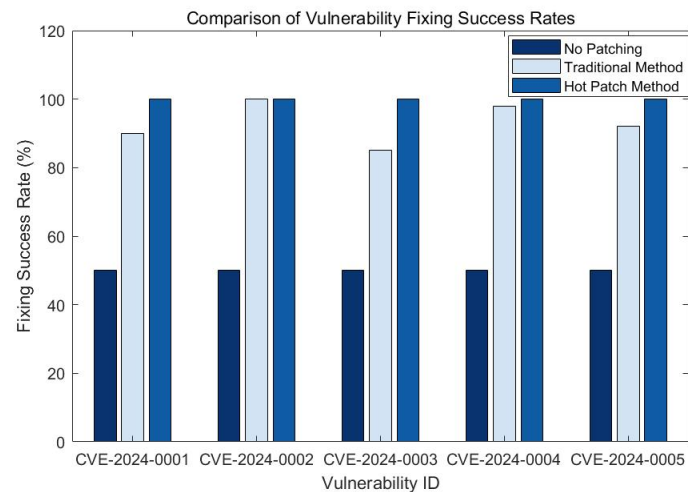


Figure 3. Comparison of Vulnerability Fixing Success Rate.

As can be seen in Figure 3, in terms of fixing success rate, the hot-patching method has a success rate of 100% on all vulnerability IDs, which is better than the traditional update method and the method without adding patch. The experimental results show that the hot-patching method can ensure a higher success rate of vulnerability fixing, which is crucial for the security and stable operation of embedded devices. Traditional update methods may not be able to completely fix the vulnerabilities in some cases, resulting in the system still having security risks. The hotfix method, on the other hand, ensures that every vulnerability is completely fixed through a dynamic repair mechanism.

4.3.3. System Stability Comparison Experiment and Result Analysis

In order to evaluate the impact of different repair methods on system stability, we compared the number of system crashes during the repair process of different methods. The experimental results are shown in Figure 4.

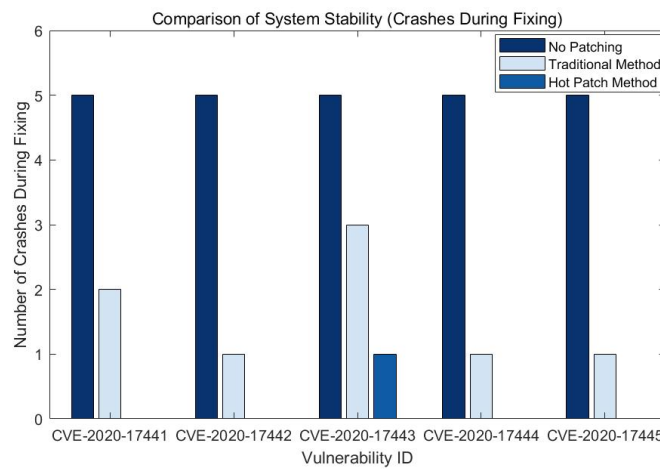


Figure 4. System Stability Comparison Chart.

As can be seen from the figure, in terms of the number of system crashes, the hot-patching approach has 0 crashes on most vulnerability IDs, which is significantly lower than the traditional update approach and the approach with no patches added. The experimental results show that the hot-patching method can significantly improve the stability of the system and reduce the number of system crashes during the repair process, thus improving the overall reliability of the system. The traditional update method may lead to system crashes and service interruptions during the repair process, whereas the hot-patching method, due to its feature of not requiring system restart, enables the system to still operate normally during the repair process, which greatly improves the stability and availability of the system.

4.3.4. System Usability Comparison Experiment and Result Analysis

In order to comprehensively evaluate the impact of different vulnerability remediation methods on system availability, we compare the system availability after remediating a vulnerability with no patch, traditional patching methods, and hotfixing methods.

$$\text{System Availability} = 100\% - (\text{Number of Crashes} \times \text{Crash Penalty}) - \left(\frac{\text{Repair Time}}{\text{Maximum Tolerance Time}} \times 10\% \right) \quad (1)$$

Assumptions: crash penalty is 2%. The maximum tolerance time is 2000 μs . The experimental results are shown in Figure 5.




Figure 5. System Stability Comparison Chart.

According to Figure 5, the hotfix method outperforms both the traditional patching method and the no-patch method in terms of system availability after remediation. For example, after the CVE-2020-17441 vulnerability was fixed, the hotfix method achieved 92% system availability, compared to 87.5% for the traditional patching method and only 60% without adding a patch. In the case of the CVE-2020-17442 vulnerability, the hotfix method had 99.1% system availability, almost a perfect score.

The results of such experiments show that the hot-patching approach is not only effective in fixing vulnerabilities, but also maximizes system availability and reduces the downtime and impact of the fixing operation. Traditional patching methods may lead to longer system downtime during the repair process, thus affecting the overall system availability. The hotfix method, however, greatly improves system stability and availability due to the fact that it does not require a system reboot, allowing the system to continue to operate normally during the remediation process.

4.3.5. Experiment and Result Analysis of Time and CPU Cycle Required for Different Patch Triggering Methods

In the experiment of different patch triggering methods, we compare the CPU cycles and time of Fixed_Patch, FPB and Debug_Monitor under different operations. The experimental results are shown in the following figures.




Figure 6. Comparison of CPU cycles required for different patch triggering methods.




Figure 7. Comparison of time required for different patch triggering methods.

From the experimental data, it can be seen that Fixed_Patch performs consistently in all cases, requires the least amount of time and CPU cycles, and is suitable for scenarios with high performance requirements. FPB mode has a higher overhead when resuming and continuing function execution (Pass), but a lower overhead when skipping or redirecting (Drop/redirect), which is suitable for scenarios with higher security requirements. Debug_Monitor has a performance between Fixed_Patch and FPB, and is suitable for scenarios that require a balance between security and performance.

Through comparative analysis, Fixed_Patch has the best performance, while FPB and Debug_Monitor provide higher flexibility and security. In practical applications, which patch triggering method to choose needs to be decided based on specific requirements and system performance indicators.

4.3.6. Patch Delay Experiments with Different Devices and Analysis of Results

In order to evaluate the performance of different devices when applying patches, we compare the patch latency of three devices, NRF52840, STM32L475 and STM32F429, under different operations. The experimental results are shown in Figure 8.




Figure 8. Patch latency on different devices.

The experimental data show that the latency of the NRF52840 is relatively stable under all test conditions, with latency fluctuating between 5.1 μ s and 5.7 μ s, demonstrating better real-time performance. The STM32L475 has a latency between 4.0 μ s and 5.0 μ s for the first five tests, with a higher latency (10 μ s) occurring in the seventh test, suggesting that in some cases performance may fluctuate. The latency of the STM32F429 is below 3 μ s in the first five tests, but increases significantly in the sixth and seventh tests to 3.5 μ s and 7 μ s, respectively.

The experimental results show that there are differences in the performance of different devices when applying patches. the NRF52840 has a more stable latency, which is suitable for application scenarios with higher real-time requirements. the STM32L475 has a lower latency, but performance fluctuations may occur in some cases, which needs to be paid attention to in practical applications. the STM32F429 has a lower latency in general but increases in longer tests, which needs to be selected according to specific needs when applying the patches. The STM32F429 has a lower latency time overall, but the latency increases in longer tests, so it should be selected according to specific needs in the application.

5. Conclusions

Analyzed by the above experimental results, the hot patching approach shows obvious advantages in several key indicators. Specifically, hot patching can significantly shorten the vulnerability fixing time while improving the stability and availability of the system. In contrast, the traditional update method may lead to long system downtime and affect normal operation, while the unapplied patches may trigger frequent system crashes and service interruptions. Overall, the hot-patching approach performs well in embedded device vulnerability remediation, effectively reducing system downtime and service interruptions, and providing strong support for system stability and availability. The performance difference between different devices when applying patches also provides an important reference for choosing appropriate remediation strategies, which helps optimize the effect of patch application and overall system performance.

Acknowledgments: This study was supported by the National Key Research and Development Program of China (Project No. 2022YFB3104300).

References

1. Niesler C, Surminski S, Davi L. HERA: Hotpatching of Embedded Real-time Applications. In: Proc. Network and Distributed Systems Security (NDSS) Symp. 2021. p. 21–25. doi: <https://doi.org/10.14722/NDSS.2021.24159>
2. Zhang Z, Wang X, Hao Q, Xu D, Wang J, Liu J, Ma J, Zhang J. Hardware-Implemented Security Processing Unit for Program Execution Monitoring and Instruction Fault Self-Repairing on Embedded Systems[J]. *Applied Sciences*, 2022, 12(7): 3584.
3. Li Z, Liu M, Wang P, Su W, Chang T, Chen X, Zhou X. Multi-ARCL: Multimodal Adaptive Relay-Based Distributed Continual Learning for Encrypted Traffic Classification[J]. *Journal of Parallel and Distributed Computing*, 2025, 201: 105083.
4. Ye H, Gu J, Martinez M, et al. Automated classification of overfitting patches with statically extracted code features. *IEEE Trans. Software Eng.* 2021;48(8):2920–2938. doi: <https://doi.org/10.1109/TSE.2021.3071750>
5. Xiong Y F, Liu X Y, Zeng M H, et al. Identifying patch correctness in test-based program repair. In: Proc. 40th Int. Conf. Software Eng. ACM: New York, 2018. p. 789–799. doi: <https://doi.org/10.1145/3180155.3180182>
6. Jiang W, Liu L. Research on hot-patching technology based on VXWORKS system. *Comput. Technol. Dev.* 2017;27(03):18–22+28.
7. Zhou M, Wang H, Li K, et al. Save the Bruised Striver: A Reliable Live Patching Framework for Protecting Real-World PLCs. In: Proc. 19th Eur. Conf. Comput. Syst. 2024. p. 1192–1207. doi: <https://doi.org/10.1145/3627703.3650068>
8. He Y, Zou Z, Sun K, Liu Z, Xu K, Wang Q, Shen C, Wang Z, Li Q. RapidPatch: Firmware Hotpatching for Real-Time Embedded Devices. In: Proc. USENIX Security Symp. 2022. p. 10–12.
9. Wang M. Exploration and application of remote hot deployment based on SpringBoot. *Inf. Comput. (Theor. Ed.)*. 2023;35(07):1–4. doi: <https://doi.org/0.3969/j.issn.1003-9767.2023.07.001>
10. Holmbacka S, et al. Lightweight Framework for Runtime Updating of C-Based Software in Embedded Systems. In: Workshop on Hot Topics in Software Upgrades. 2013.
11. Ramaswamy A, Bratus S, et al. Katana: A Hot Patching Framework for ELF Executables. In: Proc. ARES 2010. p. 507–512. doi: <https://doi.org/10.1109/ares.2010.112>
12. Payer M, et al. DynSec: On-the-fly Code Rewriting and Repair. In: Workshop on Hot Topics in Software Upgrades. 2013.
13. Chen H, Chen R-X, et al. Live updating operating systems using virtualization. In: Proc. Int. Conf. Virtual Execution Environments. 2006. doi: <https://doi.org/10.1145/1134760.1134767>
14. Mao Y, Migliore V, Nicomette V. MATANA: A Reconfigurable Framework for Runtime Attack Detection Based on the Analysis of Microarchitectural Signals[J]. *Applied Sciences*, 2022, 12(3): 1452.
15. Zhu Z, Jin Z, Wang X, Zhang J, Chen H, Choo K K R. Rethinking Transferable Adversarial Attacks with Double Adversarial Neuron Attribution[J]. *IEEE Transactions on Artificial Intelligence*, 2024.
16. Altekar G, et al. OPUS: Online Patches and Updates for Security. In: Proc. USENIX Security Symp. 2005.
17. Dong Y, et al. Towards the Detection of Inconsistencies in Public Security Vulnerability Reports. In: Proc. USENIX Security Symp. 2019.
18. Shi J, Wang Y, Su Y, et al. Analysis of MQX interrupt mechanism and interrupt program framework design based on ARM Cortex-M4. *Comput. Sci.* 2013;40(06):41–44+79. doi: <https://doi.org/10.3969/j.issn.1002-137X.2013.06.008>
19. Zhang J, Ling X, Wang Y. Design of a generic framework for assembly engineering for ARM Cortex-M series cores. *Small Microcomput. Syst.* 2021;42(11):2440–2445.
20. Zhou C. Design and Implementation of Server/Client Based Patch Management System. *Microcomput. Appl.* 2009;30(06):53–57.
21. Li Z, Wang P, Wang Z, et al. Flowganomaly: Flow-based anomaly network intrusion detection with adversarial learning. *Chin. J. Electron.* 2024;33(1):58–71.
22. Wang Z X, Li Z Y, Fu M Y, et al. Network traffic classification based on federated semi-supervised learning. *J. Syst. Archit.* 2024;149:103091.
23. Li Z, Zhang Z, Fu M, et al. A novel network flow feature scaling method based on cloud-edge collaboration. In: Proc. 2023 IEEE 22nd Int. Conf. Trust, Security and Privacy in Computing and Communications (TrustCom). IEEE: 2023. p. 1947–1953.

24. Correa R, Bermejo Higuera J R, Higuera J B, et al. Hybrid security assessment methodology for web applications[J]. *Computer Modeling in Engineering & Sciences*, 2021, 126(1): 89–124.
25. Naeem H, Alsirhani A, Alserhani F M, Ullah F, Krejcar O. Augmenting Internet of Medical Things Security: deep ensemble integration and methodological fusion[J]. *Computer Modeling in Engineering & Sciences*, 2024, 141(3): 2185–2223.
26. Zhu Z, Chen H, Zhang J, Wang X, Jin Z, Xue M, Zhu D, Choo K K R. MFABA: A More Faithful and Accelerated Boundary-Based Attribution Method for Deep Neural Networks[C]. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024, 38(15): 17228–17236.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.