computers are bad

https://computer.rip - me@computer.rip - fax: +1 (505) 926-5492

2023-11-04 nuclear safety

Nuclear weapons are complex in many ways. The basic problem of achieving criticality is difficult on its own, but deploying nuclear weapons as operational military assets involves yet more challenges. Nuclear weapons must be safe and reliable, even with the rough handling and potential of tampering and theft that are intrinsic to their military use.

Early weapon designs somewhat sidestepped the problem by being stored in inoperational condition. During the early phase of the Cold War, most weapons were "open pit" designs. Under normal conditions, the pit was stored separately from the weapon in a criticality-safe canister called a birdcage. The original three nuclear weapons stockpile sites (Manzano Base, Albuquerque NM; Killeen Base, Fort Hood TX; Clarksville Base, Fort Campbell KY) included special vaults to store the pit and assembly buildings where the pits would be installed into weapons. The pit vaults were designed not only for explosive safety but also to resist intrusion; the ability to unlock the vaults was reserved to a strictly limited number of Atomic Energy Commission personnel.

This method posed a substantial problem for nuclear deterrence, though. The process of installing the pits in the weapons was time consuming, required specially trained personnel, and wasn't particularly safe. Particularly after the dawn of ICBMs, a Soviet nuclear attack would require a rapid response, likely faster than weapons could be assembled. The problem was particularly evident when nuclear weapons were stockpiled at Strategic Air Command (SAC) bases for faster loading onto bombers. Each SAC base required a large stockpile area complete with hardened pit vaults and assembly buildings. Far more personnel had to be trained to complete the assembly process, and faster. Opportunities for mistakes that made weapons unusable, killed assembly staff, or contaminated the environment abounded.

As nuclear weapons proliferated, storing them disassembled became distinctly unsafe. It required personnel to perform sensitive operations with high explosives and radioactive materials, all under stressful conditions. It required that nuclear weapons be practical to assemble and disassemble in the field, which prevented strong anti-tampering measures.

The W-25 nuclear warhead, an approximately 220 pound, 1.7 kT weapon introduced in 1957, was the first to employ a fully sealed design. A relatively small warhead built for the Genie air-to-air missile, several thousand units would be stored fully assembled at Air Force sites. The first version of the W-25 was, by the AEC's own admission, unsafe to transport and store. It could detonate by accident, or it could be stolen.

The transition to sealed weapons changed the basic model of nuclear weapons security. Open weapons relied primarily on the pit vault, a hardened building with a bank-vault door, as the authentication mechanism. Few people had access to this vault, and two-man policies were in place and enforced by mechanical locks. Weapons stored assembled, though, lacked this degree of protection. The advent of sealed weapons presented a new possibility, though: the security measures could be installed inside of the weapon itself.

Safety elements of nuclear weapons protect against both unintentional and intentional attacks on the weapon. For example, from early on in the development of sealed implosion-type weapons "one-point safety" became common (it is now universal). One-point safe weapons have their high explosive implosion charge designed so that a detonation at any one point in the shell will never result in a nuclear yield. Instead, the imbalanced forces in the implosion assembly will tear it apart. This improper detonation produces a "fizzle yield" that will kill bystanders and scatter nuclear material, but produces orders of magnitude less explosive force and radiation dispersal than a complete nuclear detonation.

The basic concept of one-point safety is a useful example to explain the technical concepts that followed later. One-point safety is in some ways an accidental consequence of the complexity of implosion weapons: achieving a full yield requires an extremely precisely timed detonation of the entire HE shell. Weapons relied on complex (at the time) electronic firing mechanisms to achieve the required synchronization. Any failure of the firing system to produce a simultaneous detonation results in a partial yield because of the failure to achieve even implosion. One-point safety is essentially just a product of analysis (today computer modeling) to ensure that detonation of a single module of the HE shell will never result in a nuclear yield.

This one-point scenario could occur because of outside forces. For example, one-point safety is often described in terms of enemy fire. Imagine that, in combat conditions, anti-air weapons or even rifle fire strike a nuclear weapon. The shock forces will reach one side of the HE shell first. If they are sufficient to detonate it (not an easy task as very insensitive explosives are used), the one-point detonation will destroy the weapon with a fizzle yield.

We can also examine one-point safety in terms of the electrical function of the weapon. A malfunction or tampering with a weapon might cause one of the detonators to fire. The resulting one-point detonation will destroy the weapon. Achieving a nuclear yield requires that the shell be detonated in synchronization, which naturally functions as a measure of the correct operation of the firing system. Correctly firing a nuclear weapon is complex and difficult, requiring that multiple components are armed and correctly functioning. This itself serves as a safety mechanism since correct operation, difficult to achieve by intention, is unlikely to happen by accident.

Like most nuclear weapons, the W-25 received a series of modifications or "mods." The second, mod 1 (they start at 0), introduced a new safety mechanism: an environmental sensing device. The environmental sensing device allowed the weapon to fire only if certain conditions were satisfied, conditions that were indicative of the scenario the weapon was intended to fire in. The details of the ESD varied by weapon and probably even by application within a set of weapons, but the ESD generally required things like a moving a certain distance at a certain speed (determined by inertial measurements) or a certain change in altitude in order to arm the weapon. These measurements ensured that the weapon had actually been fired on a missile or dropped as a bomb before it could arm.

The environmental sensing device provides one of two basic channels of information that weapons require to arm: indication that the weapon is operating under normal conditions, like flying towards a target or falling onto one. This significantly reduces the risk of unintentional detonation.

There is a second possibility to consider, though, that of intentional detonation by an unauthorized user. A weapon could be stolen, or tampered with in place as an act of terrorism. To address this possibility, a second basic channel of input was developed: intent. For a weapon to detonate, it must be proven that an authorized user has the intent to detonate the weapon.

The implementation of these concepts has varied over time and by weapon type, but from unclassified materials a general understanding of the architecture of these safety systems can be developed. I decided to write about this topic not only because it is interesting (it certainly is), but also because many of the concepts used in the safety design of nuclear weapons are also applicable to other systems. Similar concepts are used, for example, in life-safety systems and robotics, fields where unintentional operation or tampering can cause significant harm to life and property. Some of the principles are unsurprisingly analogous to cryptographic methods used in computer security, as well.

The basic principle of weapons safety is called the strong link, weak link principle, and it is paired to the related idea of an exclusion zone. To understand this, it's helpful to remember the W-25's sealed design. For open weapons, a vault was used to store the pit. In a sealed weapon, the vault is, in a sense, built into the weapon. It's called the exclusion zone, and it can be thought of as a tamper-protected, electrically isolated chamber that contains the vital components of the weapon, including the electronic firing system.

In order to fire the weapon, the exclusion zone must be accessed, in that an electrical signal needs to be delivered to the firing system. Like the bank vaults used for pits, there is only one way into the exclusion zone, and it is tightly locked. An electrical signal must penetrate the energy barrier that surrounds the exclusion zone, and the only way to do so is by passing through a series of strong links.

The chain of events required to fire a nuclear weapon can be thought of like a physical chain used to support a load. Strong links are specifically reinforced so that they should never fail. We can also look at the design through the framework of information security, as an authentication and authorization system. Strong links are strict credential checks that will deny access under all conditions except the one in which the weapon is intended to fire: when the weapon is in suitable environmental conditions, has received an authorized intent signal, and the fuzing system calls for detonation.

One of the most important functions of the strong link is to confirm that correct environmental and intent authorization has occurred. The environmental sensing device, installed in the body of the weapon, sends its authorizing signal when its conditions are satisfied. There is some complexity here, though. One of the key concerns in weapons safety was the possibility of stray electrical signals, perhaps from static or lightning or contact with an aircraft electrical system, causing firing. The strong link needs to ensure that the authorization signal received really is from the environmental sensing device, and not a result of some electrical transient.

This verification is performed by requiring a unique signal. The unique signal is a digital message consisting of multiple bits, even when only a single bit of information (that environmental conditions are correct) needs to be conveyed. The extra bits serve only to make the message complex and unique. This way, any transient or unintentional electrical signal is extremely unlikely to match the correct pattern. We can think of this type of unique signal as an error detection mechanism, padding the message with extra bits just to verify the correctness of the important one.

Intent is a little trickier, though. It involves human input. The intent signal comes from the permissive action link, or PAL. Here, too, the concept of a unique signal is used to enable the weapon, but this time the unique signal isn't only a matter of error detection. The correct unique signal is a secret, and must be provided by a person who knows it.

Permissive action links are fascinating devices from a security perspective. The strong link is like a combination lock, and the permissive action link is the key or, more commonly, a device through which they key is entered. There have been many generations of PALs, and we are fortunate that a number of older, out of use PALs are on public display at the National Museum of Nuclear Science and History here in Albuquerque.

Here we should talk a bit about the implementation of strong links and PALs. While newer designs are likely more electronic, older designs were quite literally combination locks: electromechanical devices where a stepper motor or solenoid had to advance a clockwork mechanism in the correct pattern. It was a lot like operating a safe lock by remote. The design of PALs reflected this. Several earlier PALs are briefcases that, when opened, reveal a series of dials. An operator has to connect the PAL to the weapon, turn all the dials to the correct combination, and then press a button to send to the unique signal to the weapon.

Later PALs became very similar to the key loading devices used for military cryptography. The unique signal is programmed into volatile memory in the PAL. To arm a weapon, the PAL is connected, an operator authenticates themselves to the PAL, and then the PAL sends the stored unique signal. Like a key loader, the PAL itself incorporates measures against tampering or theft. A zeroize function is activated by tamper sensors or manually and clears the stored unique key. Too many failures by an operator to authenticate themselves also results in the stored unique signal being cleared.

Much like key loaders, PALs developed into more sophisticated devices over time with the ability to store and manage multiple unique signals, rekey weapons with new unique signals, and to authenticate the operator by more complex means. A late PAL-adjacent device on public display is the UC1583, a Compaq laptop docked to an electronic interface. This was actually a "PAL controller," meaning that it was built primarily for rekeying weapons and managing sets of keys. By this later era of nuclear weapons design, the PAL itself was typically integrated into communications systems on the delivery vehicle and provided a key to the weapon based on authorization messages received directly from military command authorities.

The next component to understand is the weak link. A strong link is intended to never fail open. A weak link is intended to easily fail closed. A very basic type of weak link would be a thermal fuse that burns out in response to high temperatures, disconnecting the firing system if the weapon is exposed to fire. In practice there can be many weak links and they serve as a protection against both accidental firing of a damaged weapon and intentional tampering. The exclusion zone design incorporates weak links such that any attempt to open the exclusion zone by force will result in weak links failing.

A special case of a weak link, or at least something that functions like a weak link, is the command disable feature on most weapons. Command disable is essentially a self-destruct capability. Details vary but, on the B61 for example, the command disable is triggered by pulling a handle that sticks out of the control panel on the side of the weapon. The command disable triggers multiple weak links, disabling various components of the weapon in hard-to-repair ways. An unauthorized user, without the expertise and resources of the weapons assembly technicians at Pantex, would find it very difficult to restore a weapon to working condition after the command disable was activated. Some weapons apparently had an explosive command disable that destroyed the firing system, but from publicly available material it seems that a more common design involved the command disable interrupting the power supply to volatile storage for unique codes and configuration information.

There are various ways to sum up these design features. First, let's revisit the overall architecture. Critical components of nuclear weapons, including both the pit itself and the electronic firing system, are contained within the exclusion zone. The exclusion zone is protected by an energy barrier that isolates it from mechanical and electrical influence. For the weapon to fire, firing signals must pass through strong links and weak links. Strong links are designed to never open without a correct unique signal, and to fail open only in extreme conditions that would have already triggered weak links. Weak links are designed to easily fail closed in abnormal situations like accidents or tampering. Both strong links and weak links can receive human input, strong links to provide intent authorization, and weak links to manually disable the weapon in a situation where custody may be lost.

The physical design of nuclear weapons is intricate and incorporates many anti-tamper and mechanical protection features, and high explosives and toxic and radioactive materials lead to hazardous working conditions. This makes the disassembly of modern nuclear weapons infamously difficult; a major challenge in the reduction of the nuclear stockpile is the backlog of weapons waiting for qualified technicians to take them apart. Command disable provides a convenience feature for this purpose, since it allows weapons to be written off the books before they can be carefully dismantled at one of very few facilities (often just one) capable of doing so. As an upside, these same properties make it difficult for an unauthorized user to circumvent the safety mechanisms in a nuclear weapon, or repair one in which weak links have failed.

Accidental arming and detonation of a nuclear weapon should not occur because the weapon will only arm on receipt of complex unique signals, including an intent signal that is secret and available only to a limited number of users (today, often only to the national command authority). Detonation of a weapon under extreme conditions like fire or mechanical shock is prevented by the denial of the strong links, the failure of the weak links, and the inherent difficulty of correctly firing a nuclear weapon. Compromise of a nuclear weapon, or detonation by an unauthorized user, is prevented by the authentication checks performed by the strong links and the tamper resistance provided by the weak links. Cryptographic features of modern PALs enhance custodial control of weapons by enabling rotation and separation of credentials.

Modern PALs particularly protect custodial control by requiring keys unknown to the personnel handling the weapons before they can be armed. These keys must be received from the national command authority as part of the order to attack, making communications infrastructure a critical part of the nuclear deterrent. It is for this reason that the United States has so many redundant, independent mechanisms of delivering attack orders, ranging from secure data networks to radio equipment on Air Force One capable of direct communication with nuclear assets.

None of this is to say that the safety and security of nuclear weapons is perfect. In fact, historical incidents suggest that nuclear weapons are sometimes surprisingly poorly protected, considering the technical measures in place. The widely reported story that the enable code for the Minuteman warhead's PAL was 00000000 is unlikely to be true as it was originally reported, but that's not to say that there are no questions about the efficacy of PAL key management. US weapons staged in other NATO countries, for example, have raised perennial concerns about effective custody of nuclear weapons and the information required to use them.

General military security incidents endanger weapons as well. Widely reported disclosures of nuclear weapon security procedures by online flash card services and even Strava do not directly compromise these on-weapon security measures but nonetheless weaken the overall, multi-layered custodial security of these weapons, making other layers more critical and more vulnerable.

Ultimately, concerns still exist about the design of the weapons themselves. Most of the US nuclear fleet is very old. Many weapons are still in service that do not incorporate the latest security precautions, and efforts to upgrade these weapons are slow and endangered by many programmatic problems. Only in 1987 was the entire arsenal equipped with PALs, and in 2004 all weapons were equipped with cryptographic rekeying capability.

PALs, or something like them, are becoming the international norm. The Soviet Union developed similar security systems for their weapons, and allies of the United States often use US-designed PALs or similar under technology sharing agreements. Pakistan, though, remains a notable exception. There are still weapons in service in various parts of the world without this type of protection. Efforts to improve that situation are politically complex and run into many of the same challenges as counterproliferation in general.

Nuclear weapons are perhaps safer than you think, but that's certainly not to say that they are safe.

[1] This ”popular fact” comes from an account by a single former missileer. Based on statements by other missile officers and from the Air Force itself, the reality seems to be complex. The 00000000 code may have been used before the locking mechanism was officially placed in service, during a transitional stage when technical safeguards had just been installed but missile crews were still operating on procedures developed before their introduction. Once the locking mechanism was placed in service and missile crews were permitted to deviate from the former strict two-man policy, ”real” randomized secret codes were used.