On the Safest Way to Kill

by John Ridgway

The fascination of hunting as a sport depends almost wholly on whether you are at the right or wrong end of the gun.

P. G. Wodehouse, *The Adventures of Sally*

I like to think that business is very much like hunting. Basically, people hotly pursue their goals, and there are those who will benefit from the pursuit and others who will… well, maybe not benefit so much. To capture this idea, organisations like to talk of differing stakeholders, which is just a lofty way of saying that not everyone shares the same goals and concerns.

Stakeholding can make the construction of a safety case very interesting (if not to say contentious) since one person's safety can often set the scene for another's downfall. Nowhere is this contrast as stark as it is for nuclear weaponry. When nuclear weapons work safely, people are still allowed to die in their millions. Yet an unsafe weapon might only kill a handful of people, albeit an unintended handful. For example, a maintenance team could be killed by faulty detonators. Worse still, your warhead may fail to go off altogether, thereby allowing the enemy a greater opportunity to kill millions. And, of course, we are talking about the 'wrong' millions here.

Military analysts refer to the concept of Always/Never, which means that a weapon must always fire when required but never fire when not required. Consequently, if safety devices are to be applied, they cannot be allowed to compromise the mission – the weapon bearer's stakeholding must be protected under all circumstances. This uneasy balance between the demands of both mission and safety is central to the development of the world's nuclear arsenal, and it is a good reason why you shouldn't sleep too easily tonight.

How Safe Do Nuclear Weapons Have to Be?

Before I question too deeply the logic of developing a safety case that permits mass destruction, I need to consider the level of risk that can be tolerated on behalf of the population you claim to be protecting.

The first attempt to answer this question was made by the United States Army in 1955. They calculated how many fatalities had resulted from floods, earthquakes and other natural disasters in the preceding fifty years and used this as the target mortality rate for nuclear accidents. From this they calculated that the acceptable odds for an atomic (fission) weapon to explode accidentally on U.S. soil should be 1 in 250 during the course of a year. The same target equated to a 1 in 100,000 chance for thermonuclear weapons. Remember, we are talking about accidental death here. Determining the acceptable number of communist deaths required quite a different calculation.

In 1957 an organisation known as the Armed Forces Special Weapons Project revisited the question and concluded that there should be no more than a one in five chance of an accidental thermonuclear detonation within a given decade. The same organisation condoned the virtual certainty of a fission weapon accidentally detonating within the same timescale. This was the military viewpoint. One can imagine that civilians would come up with a different answer.

Since those early days, nuclear weaponry has seen many changes. Furthermore, the social and political backdrops have

evolved. Consequently, the safety standards required for nuclear weapons have undergone significant reassessment. The latest statement on the subject (chapter 2 of US Defense Standard DoD 31 50.2-M, 1996) states that the probability that a nuclear warhead might accidentally detonate, whilst in normal storage and operational conditions, shall not exceed one in a billion per warhead lifetime. However, this target drops to one in a million for abnormal environments such as those resulting from fires. Keep in mind that this limit is per warhead, so a goodly stockpile will increase the risk significantly. It seems, even now, no one expects nuclear weapons to be entirely safe, and many expert authorities fear that an accidental or unauthorised detonation is only a matter of time.

So How Safe Are They?

Or, more to the point, should one really expect a straight answer to this question? For example, when it was proposed that Strategic Air Command bombers should carry live nuclear bombs as part of a permanent airborne readiness, President Eisenhower was advised that the chances that any aviation accident could result in a nuclear explosion were, "essentially zero". In truth, zero was the only probability that could be confidently ruled out. And given that B52 bombers were, at that time, crashing once in about every 20,000 flying hours, the real prospects of avoiding a very bad day at the Oval Office were not nearly as good as Mr President thought.

So, I won't insult your intelligence by plying you with 'official' facts and figures. Nevertheless, a survey of things that might go wrong can still be illuminating.[1] Firstly, I'll address questions of physical vulnerability, then I will take a quick look at failures of command and control. As I do so, I will discuss safeguards and the ever-present conflict between the pursuit of both safety and mission objectives.

Wooden Bombs

In the earliest days of atomic weapons, the risk that a bomb might accidentally detonate was minimised by delaying final assembly until the point of deployment.[2] Prior to that point, the fissile component would be stored separately from the high-explosive components used to detonate it. Whilst this arrangement improved the safety margins whilst in storage, it was bad news as far as operational readiness was concerned. Furthermore, the weapons were difficult to maintain and very bulky and heavy. Weapons development therefore focused on creating relatively lightweight[3], preassumbled bombs that could be safely stored at dispersed sites and then loaded and primed within minutes – so-called "wooden bombs", sat inertly on the shelf awaiting use. Since such weapons would have their fissile core confined within a solid metal casing, they were said to have a sealed pit.

One-Point Safety

The resulting weapons were smaller and more powerful but had the disadvantage of being ready to bang if subjected to physical trauma, as in a bullet strike or a fire. Consequently, they were required to be "one-point safe". This decreed that detonation at a single point within a bomb's high explosive system (used to implode the fissile core) shall not result in a nuclear explosion.[4] Furthermore, any non-nuclear explosion's yield should not exceed the equivalent of four pounds of TNT (in consideration of the damage that a non-nuclear explosion can still cause on a ship). Notwithstanding these arrangements, the behaviour of high explosives under abnormal conditions, such as in a fuel fire, is sufficiently uncertain as to warrant extreme caution when dealing with such conditions. Each make of bomb had a "time factor", typically only a few minutes, after which every

firefighter and their dog was advised to run to the hills.

More Things to Wory About

Irrespective of one-point safety, two serious concerns remained to preoccupy the safety conscious combatant:

Firstly, a non-nuclear explosion may still result in the dispersal of highly radioactive material. And if this happens to be plutonium, you are dealing with a very serious risk indeed, since plutonium-contaminated dust can be easily inhaled, to deadly effect. Worse still, plutonium has a half-life of about 24,000 years, so any contamination is likely to outstay its welcome.

Secondly, early sealed pit weapons ran the risk of accidental detonation caused by a rogue firing-signal. This could result from a maintenance accident, a lightning strike or a fire. To counter this risk, a trajectory sensing switch was developed, whereby the firing mechanism would only become operational following weapons launch or release. Prior to that point, it should be physically impossible for current to flow to the detonators. But Mother Nature is ingenious, so no mechanism is guaranteed to carry zero risk.[5] And the harder you try to reduce the risk of accidental fusing, the more likely you are to end up with a dud when you press the button. The military have good reason to worry about this. A safety device added to Polaris missiles was later found to be subject to corrosion, rendering the missile dud. In fact, a 1963 routine maintenance check discovered that at least 75% of the Polaris fleet's missiles were destined to disappoint.

If in Doubt, Build More Bombs

Perversely, the early development of safety mechanisms was to make the world less safe. Concerned that a good many of their weapons might now be of the "Never/Never" persuasion, the military compensated by increasing the nuclear stockpile, with the intention of hitting any given target with many and varied warheads. This clearly increases the risk that a full-scale nuclear conflict might be instigated in error, making that of a single, accidental detonation appear quite attractive in comparison. So this is probably a good time to discuss how a nuclear contretemps could happen and how it may be avoided.

The Importance of Self-Control

First and foremost, protection[6] from nuclear attack requires that one



possesses an early warning system that is both effective and trustworthy. Consequently, any indication of an incoming attack must be properly evaluated before a response may be decided. However, the checks and balances required to avoid responding to a false alarm inevitably take up valuable time, and the "use it or lose it" principle is bound to influence deliberations. Clearly, therefore, the risk of hastily responding to a false alarm has to be suitably balanced with that of failing to respond in a timely fashion. The good news is that the military are able to practise their response using computer-simulated attack and so, presumably, have tested their procedures to remove any inefficiencies. The

not so good news is that in 1979 someone in the North American Aerospace Defense Command loaded a training tape onto a computer without anyone realising and, in so doing, very nearly precipitated World War Three. Catastrophe was averted only because surveillance satellites failed to provide corroborating evidence of an attack. Other close calls have resulted from mistaking a scientific rocket for an ICBM, a faulty computer chip, sunlight reflected from high-altitude clouds, and mistaking the moon for an incoming salvo of ICBMs.

As a further concern, one of your nuclear weapons may fall into the wrong hands, and there is always the chance that one of your own personnel may not be on message. Consequently, controls and arrangements that help avoid an unauthorised launch[7] are essential. These include: Permissive Action Links (PALs), i.e. devices that require a coded input to enable a device's firing mechanism; centrally-controlled code management; dual key arrangements, with mutual monitoring of potentially rogue behaviour; secure and reliable communications for authorisation and confirmatory dialogues; effective security protection for arsenals and armed delivery systems; and appropriate psychological profiling and selection of personnel. I could go on, but I doubt if I could convince you that the arrangements amount to a fool-proof system. As long as battle readiness remains a priority, the Devil will always find a way.[8]

From Stakeholder Management to Holocaust

The development of nuclear weaponry provides a prime example of the efforts required to meet the Always/Never challenge. However, war is what happens when stakeholder management gets out of hand, so it is no surprise that, in the battlefield, Always/Never makes sense only when seen from the appropriate stakeholder perspective. Not only are safety and mission efficiency arrangements often at odds with each other, a successful balance between them is likely to maximise the death toll, once the slaughter of all stakeholders has been taken into account. Safety arguments can look very odd when reliability equates to lethality, since the concepts of the right people to kill and the right circumstances in which to kill have a very shaky ethical foundation.

Finally, it's not as if stakeholder conflict is unique to military systems. Even in civilian applications, your stakeholder management obligations may lead you to make some morally dubious decisions. For example, what price public safety if it requires that the lives of the emergency services be placed at risk? The bottom line is that if you ever find yourself constructing a safety argument and asking yourself, "What would Jesus do?", you might want to consider getting another job.

Notes:

[1] In this article I do not attempt a quantitative assessment of the safety risk. However, those who are interested may wish to consult On the Risk of an Accidental or Unauthorised Nuclear Detonation – RM-2251 U.S. Air Force Project Rand (1958), in which Bayesian techniques are used to assess probabilities based upon "accident opportunities" and their incident rates.

[2] So, for example, an airborne weapon's nuclear core would only be inserted after take-off. To further reduce the risk, nuclear missions were not to be flown from U.S. bases. Instead, expendable airfields such as those in Norfolk, England were to be used. As far as the U.S. government was concerned, the good people of Norwich were the wrong kind of stakeholders.

[3] The weight and size reduction was principally achieved by using solid electrolytes to provide the power to activate the warhead, rather than relatively heavy, liquid electrolytes. In addition, a tritium and deuterium gas was used to boost the weapon's yield, whilst enabling the reduction of fissile material. Note that the fission boost is provided by extra neutrons produced following fusion of the lithium and deuterium. However, the energy released by the fusion

does not itself significantly contribute directly to the yield, so the weapons are still basically fission devices.

[4] One way to ensure this is to design the bomb's core so that, notwithstanding its implosion, reaching criticality still requires the boost neutrons that are released following fusion of the surrounding lithium and deuterium gas. The mechanism that activates this fusion is independent of that which detonates the high explosives. Consequently, an accidental activation of the high explosive implosion system will always be insufficient, on its own, to instigate a nuclear reaction.

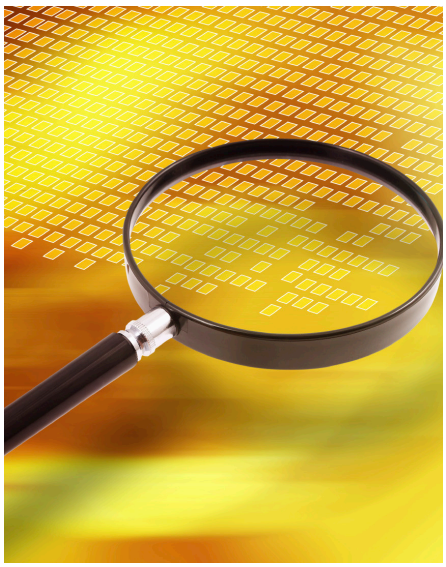
[5] For example, short circuits can result from melted solder, charred plastic or from arcing between newly proximate wiring. Errant currents might then energise the firing mechanism. Such risks are minimised by isolating the firing mechanism within a physical barrier and then using so-called 'strong links' to connect to external components. 'Weak links' within the isolated section are then designed to predictably fail 'open' in abnormal environmental conditions, long before the strong links have chance to fail in a potentially unpredictable fashion.

[6] By 'protection', I mean ensuring that your nation's imminent destruction shall not go unnoticed.

[7]The U.S. terminology for nuclear mishaps is provided by DoD Directive 5230.16. Apparently, an unauthorised launch of an ICBM, resulting in detonation, is an example of a so-called 'NUCFLASH'. I wonder how long it took them to come up with that one.

[8] To expedite launch procedures, the security code required to activate the Minuteman missile was set to '00000000' in all silos.

John Ridgway *is a retired analyst who spent most of his career working for Serco Transportation Systems. Having failed to change the world, he now lives in quiet retirement with his wife and two dogs. He can be contacted at j.ridgway4@ntlworld.com*



Requirements in the Loop: The Future of Model-Based System Engineering

by Yves Génevaux, Ken Nathan, Fabien Gaucher & Bruno Joffret

In this article, we consider the level of maturity of Systems Engineering as compared with other phrases of embedded systems development. Although much progress has been made over the years with the introduction of Requirements Management tools, Traceability tools and support for organising the elaboration of an architecture with Model-Based System Engineering (MBSE), essentially through various SysML flavours, there still remains a vast area for improvement. It is known that something of the order of 70% of the errors introduced during the various activities of an embedded system development project are actually introduced during the specification phase, whereas only some 4% of the errors are discovered during this phase. Figure 1 illustrates the percentage of errors typically introduced and detected during the development process. Given that the general consensus is that the earlier you find the errors, the more efficient and cost-effective it is, it is clear that the processes and tools used during the specification phase need to be improved.

One blocking factor to improvement is that most of the information relating to the expected behaviour of the system and of its components remains in a number of requirements written in natural language. Imagine dozens, or hundreds, of requirements that are to be considered in parallel, where any of them can be 'activated' at any