Technology
Innovation
Institute

A ZERO-TRUST APPROACH TO AUTONOMIC SWARM SECURITY, RESILIENCE, AND SAFETY

What Swarm Safety Could Learn From Zero-Trust Security

Table of Contents

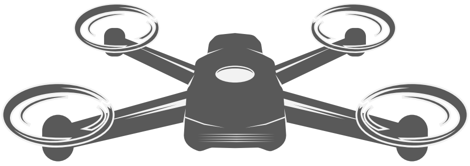

Extending Safety Assurance From Individual Robots To Swarms

Since the turn of the century, research engineers have published thousands of papers on improving the security, resilience, and safety of the individual autonomous “things” that constitute the Internet of Things—devices like automobiles, aerial drones, robots, and countless other kinds of equipment. Certainly, these unmanned vehicles—UAVs in the air, UGVs on the ground, and USVs on the sea—must safely negotiate the journey from start to destination. But they also face the vastly more complex problem of sharing the sky, roads, and sea-lanes with other traffic, both manned and unmanned. Thus, governments and commercial consortia are laying the groundwork for smart systems for air traffic control, automated road transport management, and other systems for sharing travel space safely.

The technologies have come a long way, and the pace is accelerating; however, more research is needed before collections of autonomous vehicles can be deemed workable, secure, robust, and safe. For example, an autonomous chaos was created in Austin, Texas indicating a need for an orchestration layer to manage autonomous systems [21]. Much of the recent work on UAV/UGV/USV/drone swarms since 2000 focuses on emergency applications: spotting fires, monitoring floods, assisting rescue workers in natural disasters, and the like, particularly in isolated areas. In these situations, teams of people and swarms of drones must come together quickly and work closely in harsh conditions under drastic time constraints. The lessons learned in emergencies can then often be applied to more routine applications in ground transportation, air delivery, manufacturing, warehouse automation, construction, and agriculture, and other domains.

In all cases, comprehensive systems-level approaches will be required to extend the safety engineering of individual autonomous devices in to complex-system realm of swarms, and even swarms of swarms, where unanticipated situations and emergent behaviors can arise.

This is autonomic computing: the system automatically manages itself. Like the autonomic nervous system—which regulates involuntary responses like heartbeat, blood pressure, and breathing—autonomic computing is a distributed system that manages machines’ responses to complex and unpredictable situations. Harel et al. argue that “an autonomics foundation will eventually lead to trustworthy hardware/software systems.”[1] They call for work on three main challenges: 1) Specifying autonomous behavior in the face of unpredictability; 2) carrying out faithful analysis of system behavior in a rich environment; and 3) building such systems by combining executable models using traditional software engineering, AI, and machine learning.

On its own, an autonomous system managing a swarm of drones (for example, Ghaf platform for edge-device management developed at TII’s Secure Systems Research Center; the Ghaf tree is highly resilient and flourishes even in harsh deserts) can safely respond to new situations. [2] Autonomics provides a framework for developing systems that can safely manage themselves.

An autonomic swarm leverages collective intelligence to ensure the safety of individual machine and human elements working together. SSRC’s goal is to make autonomous systems into autonomic systems with humans in the decision-making loop (“humans-on-the-loop”), rather than “humans-in-the-loop,” with continuous operational involvement.

The Harel-Marron-Sifakis principles can help create autonomic swarm safety mechanisms, including multiple interconnected feedback loops across devices, networks, and control algorithms. This will require multidisciplinary approaches that integrate robotics, artificial intelligence, networking, and systems engineering. By leveraging advances in these fields, developers may create drone swarm systems that can operate safely, efficiently, and effectively in a wide range of applications.

Autonomic swarm safety could also draw on zero-trust security principles: constant vigilance, taking external and internal threats as a given, authenticating every contact. These measures should be deployed to protect large, decentralized systems across every level of software, hardware, and system stack. Zero-trust security on its own can reduce the "blast radius" of cyber-attacks, while also reducing the impacts of faulty sensors, control systems, communication systems, and AI algorithms.

High security is the foundation of safety and resiliency. This will provide a base layer for building secure, resilient, and safe swarms to harden the sense-decide-act loop at the heart of autonomous systems. Elements can include hardened operating systems like seL4 and robust reduced-instruction-set CPUs, which reduce the number of potential points of attack. (Both seL4 and RISC-V systems are under study and development at TII. [3][4])

Overall, zero trust principles play critical roles in extending autonomic feedback loops to drone swarms by ensuring that each drone in the swarm can be trusted to act in a secure, reliable, accountable, and flexible manner. Such drone swarms can achieve their objectives effectively and safely while minimizing the risk of cyber-attacks, accidents, and other unintended consequences.

While early research focuses on drone swarms, what we learn could eventually improve the safety of autonomous infrastructures for controlling buildings, factory equipment, and smart cities.

Essential Building Blocks of Autonomics Include

Validate - Applying Zero-Trust principles to provide a base layer of integrity across all autonomic processes: sensing, decision-making, acting, and networking.

Sense - Resilient sensing by fusing inputs from multiple sensors (and multiple types of sensors) to mitigate the impact of faulty sensors and poor interpretation.

Decide - Collective resilient decision-making.

Act - Extending AI alignment to swarms to ensure they behave as intended.

Network - Connect/ collaborate via resilient mesh.

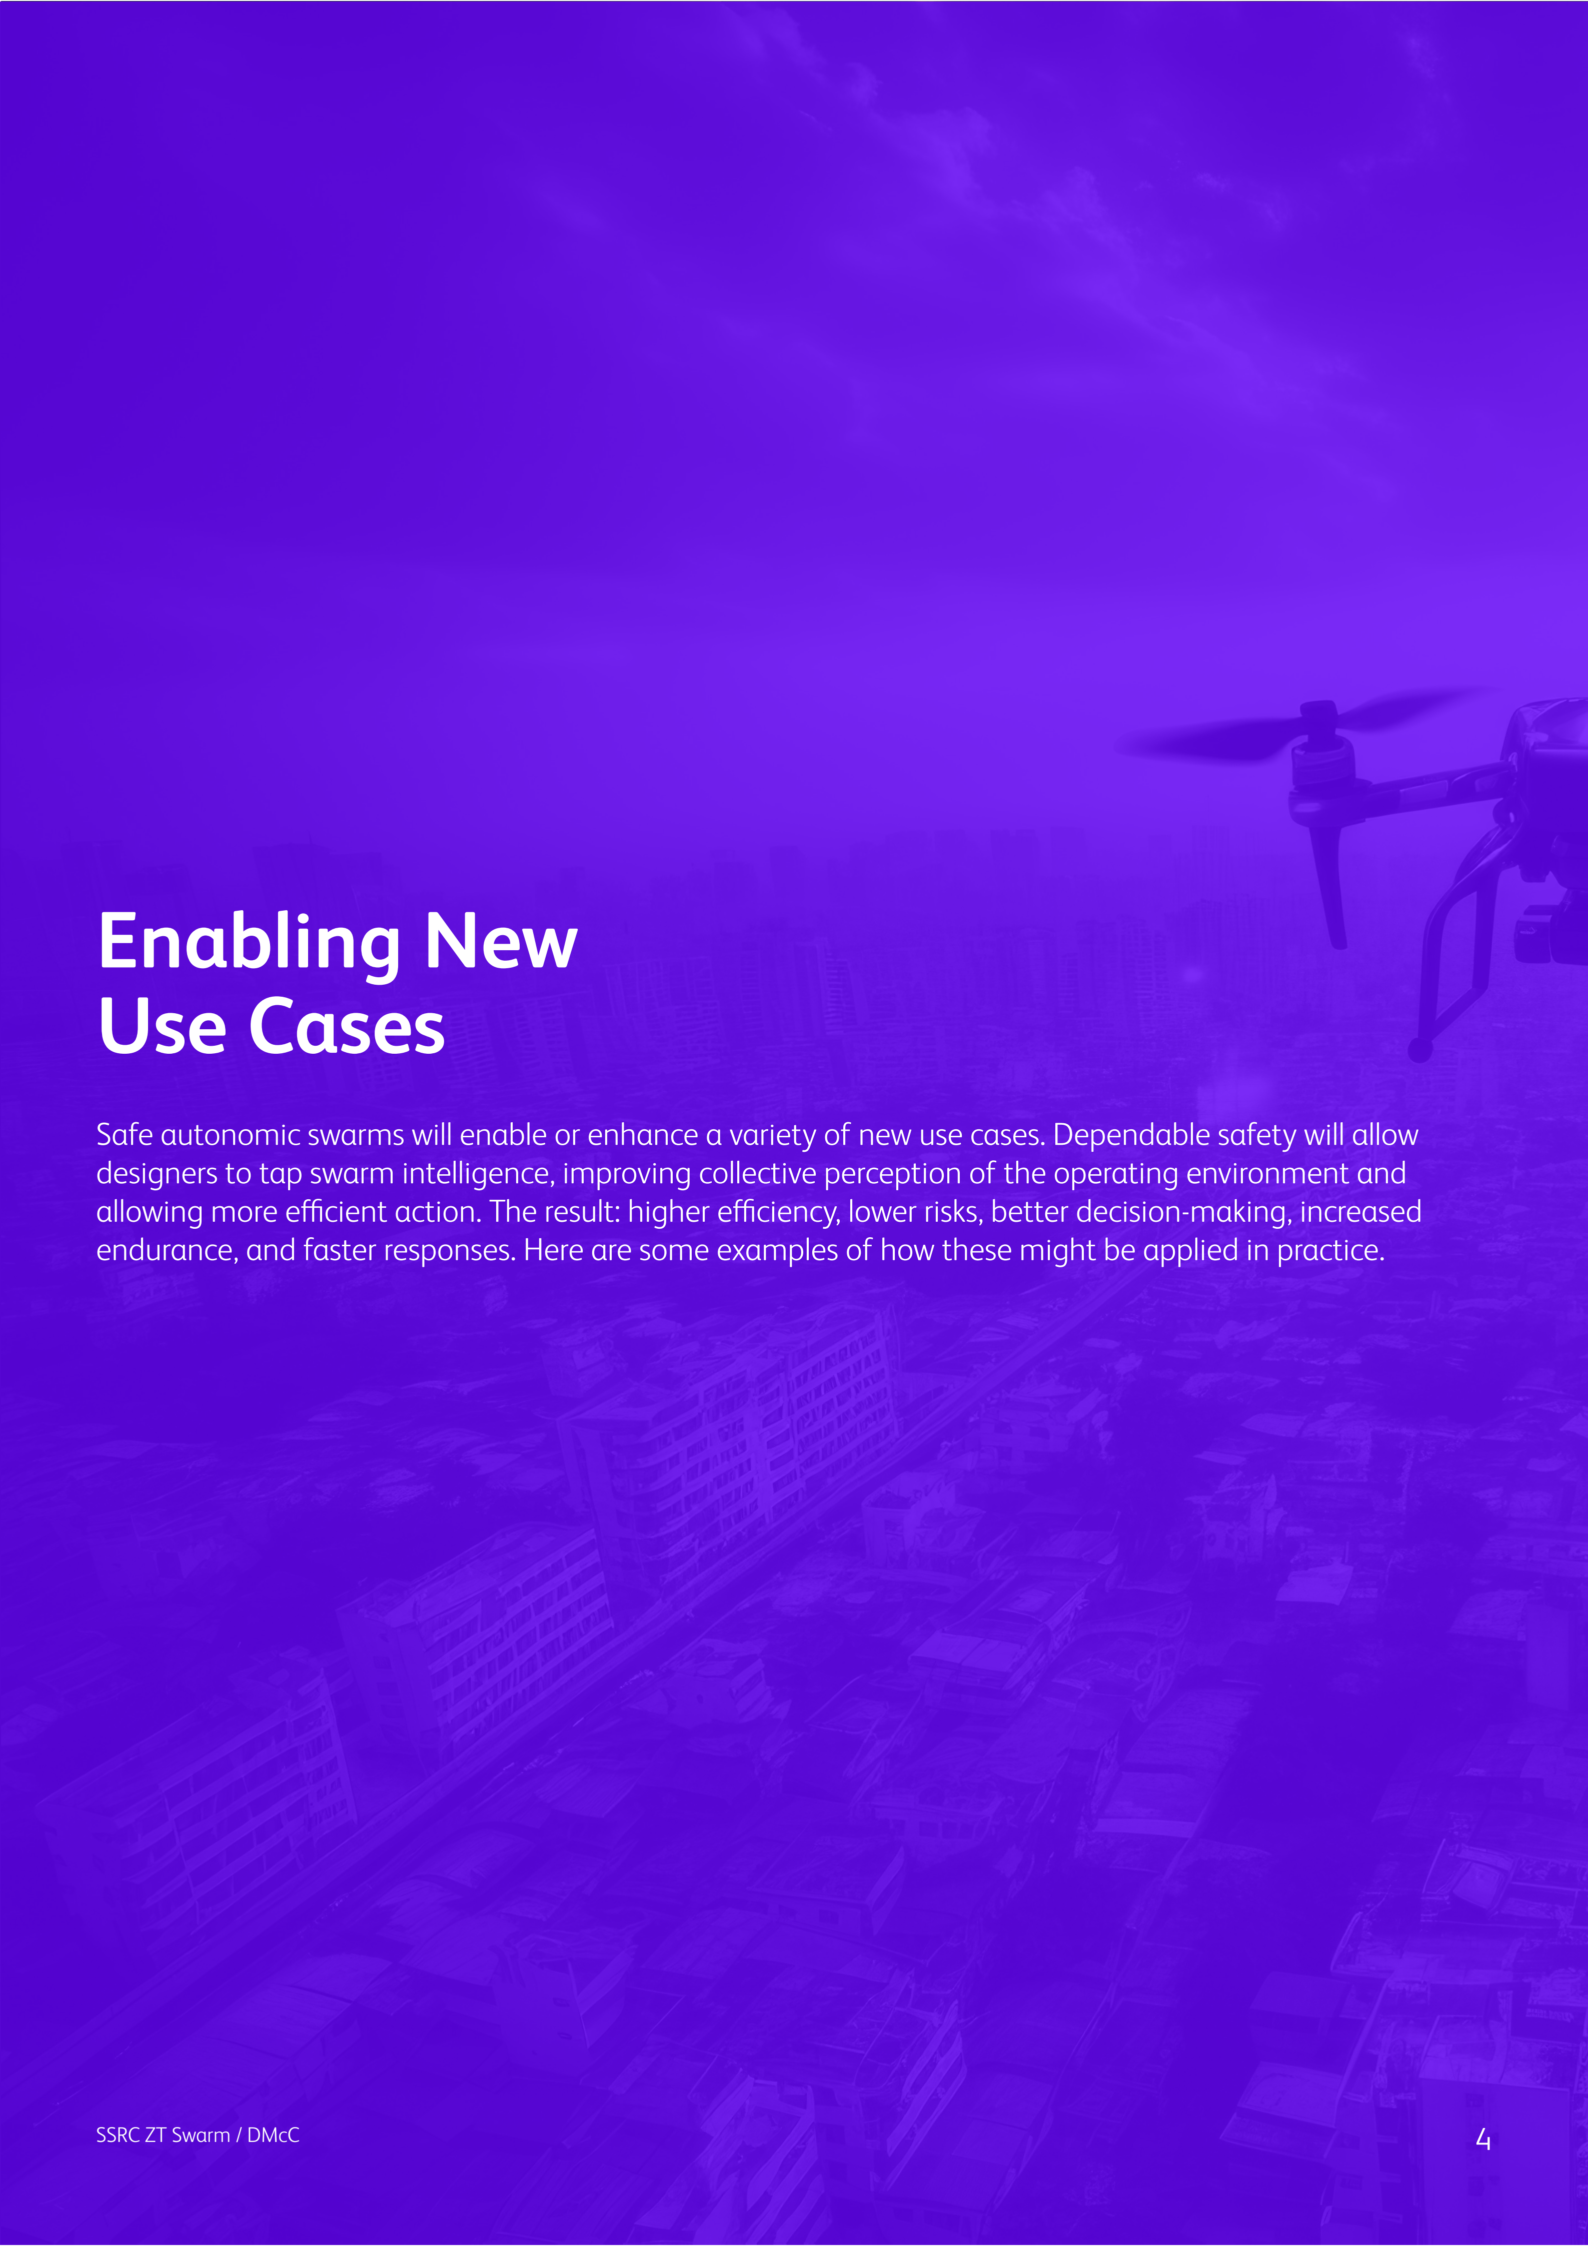

Enabling New Use Cases

Safe autonomic swarms will enable or enhance a variety of new use cases. Dependable safety will allow designers to tap swarm intelligence, improving collective perception of the operating environment and allowing more efficient action. The result: higher efficiency, lower risks, better decision-making, increased endurance, and faster responses. Here are some examples of how these might be applied in practice.

Firefighting

Emergencies like forest fires or wildfires in remote areas demand utmost performance from workers on the ground and their equipment. Fires move and change quickly, threatening life and property over large areas—often areas where roads and communication are both poor. The smoke, roaring flames, and winds can make it hard for men and women working near the blaze to see and hear.

Emergency drone systems can work closely with firefighters, uniting elements from central command and control to individuals working at the fire-head.

A swarm of aerial drones (UAVs) and unmanned ground vehicles (UGVs) could be deployed to collectively assess the exact size and shape of a large fire beyond the reach of existing wireless networks. Autonomous control centers and air and ground "fog" communications relays communicate with edge UAVs and UGVs to relay information and instructions to firefighters on the ground.

Firefighters report back on local situations, while sensors in their clothes and equipment relay images, sound, and the first responders' vital signs. The autonomous drone system combines this data with observations from edge and fog units to build a comprehensive picture of the fire. Field commanders use the data to allocate resources, and the firefighters on the ground get early warnings of impending danger.

An autonomous swarm can "self-heal" to reconfigure the coverage pattern if a unit is lost or the fire suddenly changes course. The collective data could be orchestrated into a digital twin, a precise digital representation of the geography and state of the fire that commanders in the field can used to run "what if" evaluations of multiple fire-fighting strategies to reduce overall damage.

SCMD: FIRST RESPONDER USE CASE

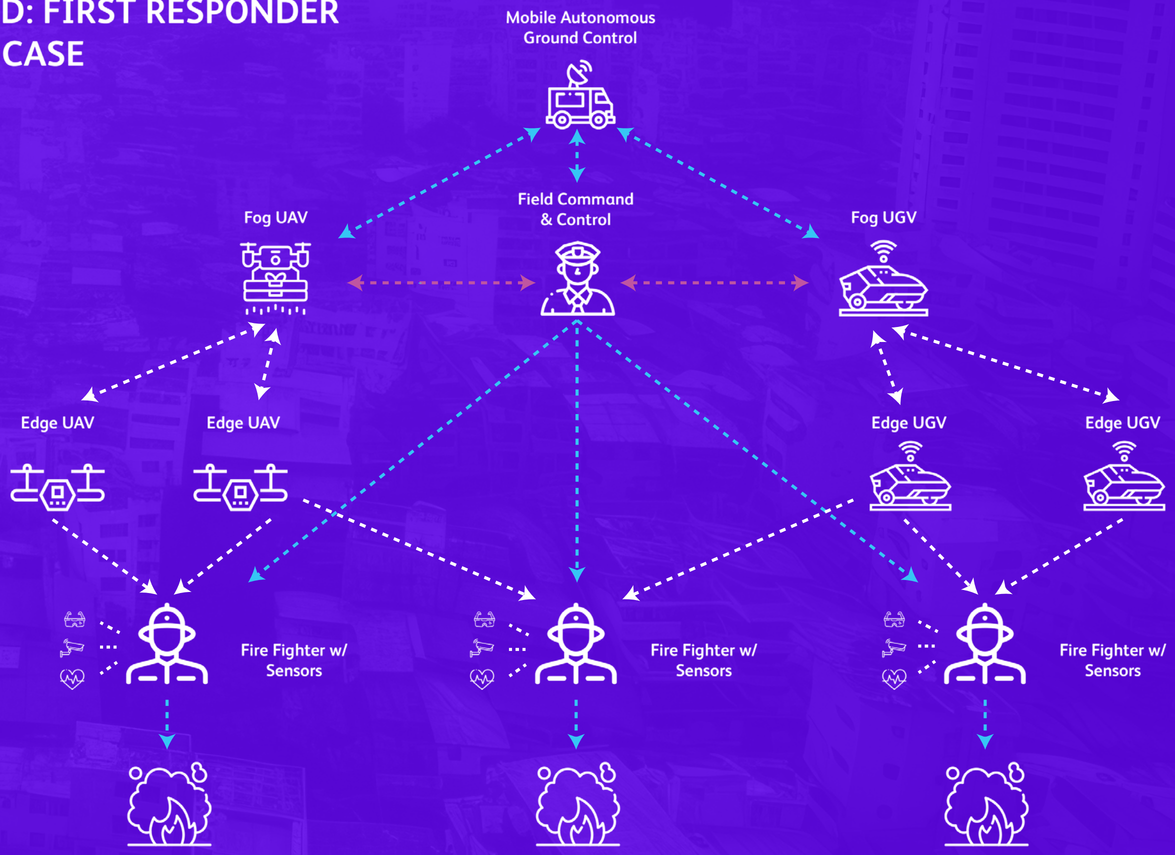

Wildfires and forest fires in remote areas can be dangerous, fast-developing, and chaotic. There may be no communication infrastructure. Visibility and hearing may be sharply curtailed. A tiered command-and-control system links human supervisors and firefighters (equipped with sensors and communications equipment) via a network of central control, "fog" air and ground units that extend communications, and "edge" units that communicate directly with fire crews and track location of the blaze.

Search And Rescue

Coordinated aerial drone swarms could reduce coverage overlaps to span wider areas to speed up searches. They can gather vital information about factors like the number of missing persons, the characteristics of the terrain, and the nature and locations of potential dangers. Mesh networks would allow them to stay connected with one another out of range of their base stations—underground, inside damaged buildings, or where existing wireless infrastructure has been destroyed.

Art And Entertainment

Drone swarms today are used to produce aerial light shows akin to fireworks. Autonomic safety assurance would allow these to operate closer to humans for more engaging (and spectacular) experiences.

Environmental Monitoring

Aerial and underwater drone swarms equipped with environmental monitoring sensors can gather data on air quality, water quality, and animal behavior. They detect anomalous conditions and allot more analytical resources where needed.

Intelligent Farming

A swarm of drones and autonomous tractors can monitor crop health, improve irrigation, and apply fertilizers or pesticides with increased precision. They can model field geography to anticipate water run-off or track key indicators like soil water content, pH, or nutrient levels, calling the famer’s attention to problem areas while improving yields and reducing use of fertilizer, herbicides, and pesticides.

Warehouse Management

Centrally controlled drones are used today to move materials, inventory products, and pick items for shipping—often in areas that are off-limits to humans for safety reasons. Better autonomous safety could improve robot-human collaboration, or increase operational flexibility—even allowing the system to reconfigure the warehouse layout on the fly, in response to changing demand patterns.

Infrastructure Inspection

Inspectors already use drones to examine infrastructure like bridges, pipelines, and power lines. By becoming even safter, drones would allow engineers to build near-real-time digital twins to pinpoint structural problems, model repair strategies, and fix problems faster and more cost-effectively than ever before.

Adaptive, Decentralized Building Controls

Autonomic principles could be extended to adaptive, decentralized building control. Collaborative cyber-physical control networks could help electrical, HVAC, water, fire, and communication systems adapt quickly to changing routine or emergency conditions, and balancing these responses with the conditions of neighboring building zones—maintaining temperatures when a heating unit goes down, for example.

Current Standards for Individual Drones

Researchers and engineers have made considerable progress in improving the safety of individual autonomous cars, UAVs, and IoT devices. The ISO 26262 standard focuses on the functional safety of electrical and electronic systems in road vehicles.[5] Meanwhile, the ISO/PAS 21448 SOTIF (Safety of the Intended Functionality) standard addresses hazards that can arise from the system's intended functionality, even when all systems function correctly. It focuses on identifying and managing potential hazards arising from limitations or insufficiencies in the system's design and intended use cases. The Society of Automotive Engineers' SAE J3016 standard nomenclature includes a schema for classifying the various levels of autonomous capability ranging from levels 0 to 5 that is used for certifying the use of autonomous vehicles on public roads.

Several frameworks have similarly been designed for unmanned aerial vehicles (UAVs). ASTM F38 covers aspects like airworthiness, systems performance, flight operations and safety.[6] ISO 21384-3 covers such aspects as identification of potential hazards, risk assessment, operational procedures, system reliability and maintenance.[7] The European Union Aviation Safety Agency (EASA) has developed a risk assessment methodology called Specific Operations Risk Assessment (SORA) for assessing the operational safety of specific drone applications.[8]

In the realm of IoT devices, several standards and frameworks have been developed to protect the security of IoT devices. IEC 62443 focuses on the security of industrial automation and control systems.[9] The Cloud Security Alliance (CSA) CSA IoT Controls Framework provides guidelines and best practices for securing IoT devices, networks, and data, covering various domains such as identity management, secure development practices, and security monitoring.[10] The European Union Agency for Cybersecurity (ENISA) has published Baseline Security Recommendations for IoT, a set of security recommendations for IoT devices and systems.[11]

There are also several efforts to provide large-scale UAV air-traffic-control systems that allow UAVs to operate beyond the visual line of sight (BVLOS). These efforts provide the infrastructure for safely controlling independent fleets of UAVs in concert with a centralized control system. The US is working on Unmanned Aircraft System Traffic Management (UTM) to develop services, roles and responsibilities, information architecture, data exchange protocols, software functions, infrastructure, and performance requirements for enabling the management of low-altitude uncontrolled drone operations.[12]

EASA is developing a U-space (think of it as airspace for unmanned aerial systems) regulatory framework in Europe to guide the development of similar UTM infrastructure under the auspices of Regulation (EU) 2021/664 that addresses these systems' technical and operational characteristics.[8]

TII

Technology
Innovation
Institute

Safety Assurance for Emergent Behavior

These efforts will all play an important role in the widespread adoption of drones. It is also important, however, to consider safety in the emergent behavior of interacting autonomous swarms, a relatively new realm of engineering. The University of Bristol has developed a framework called AERoS (Assurance of Emergent Behavior in Autonomous Robotic Swarms) developed from the Assurance of Machine Learning in Autonomous Systems.[13]

The University of Bristol researchers observed that existing safety standards focus on assuring safety of robots at the individual level but need to cover the safety implications of the emergent behavior of swarms. This framework is a good first effort in thinking about how to simulate and address new problems that can arise at the system-of-systems level. The Bristol group also calls for more work on the trustworthiness of drone systems—more work on ethics, governance, and regulation of autonomous system design and operation. One promising direction is to explore how Carl Macrae's Structural, Organizational, Technological, Epistemic, and Cultural (SOTEC) framework could help identify and mitigate new sources of risks in autonomous systems.[14] This would allow autonomous swarms to operate within a wider operational context.

Some principles developed for autonomic computing, autonomic networks, and zero-trust security could play important roles in improving security, safety, and resilience at a swarm level. Autonomic computing provides a way of creating failover strategies when individual processors or drones fail or misbehave. Similarly, autonomic networks would use secure wireless meshes that can revert to backup mechanisms such as visible light communication when radios go offline or are jammed. Secure trust methodologies provide a way of thinking about the security requirements and the impacts of faulty sensors or actuators.

Autonomous vehicles sometimes suffer from "hallucinations," mistaking anomalous sensor data for hazards or obstacles, which can result in problems like phantom braking. In one particularly alarming event, in November 2022, a Tesla suddenly braked in heavy San Francisco traffic, causing an 8-car pileup, 18 injuries, and stalling traffic for hours.[15]

It's also important to consider the safety hazards of malicious attacks on drone swarms. In 2021 the US Federal Aviation Administration recorded more than 9,700 laser attacks on airplanes. New RF hacking gear inspired by the Flipper Zero, which allows people to imitate remote controls, might give unruly teenagers or criminals unprecedented new ways to vandalize drone swarms. These are issues with individual robots as well, of course, but they can become exponentially more complex with swarms of collaborative robots or drones.

Building on a base zero-trust framework could mitigate the impact of these kinds of new attacks. A drone swarm needs a way to adapt to errant misperceptions of the environment; otherwise, a failure mode can cascade across the swarm, creating bigger problems. A zero-trust framework provides a way to discount input from an attacked robot or sensor before the harm spreads. The other robots can make safety adjustments when they determine that one of them is getting faulty data or not operating as intended.

A Framework For Autonomic Drone Safety

Several elements must be considered to extend autonomic feedback loops to ensure efficient, effective, and safe drone swarms. This work needs to build on the existing body of research, engineering, and safety standards for ensuring the safety of individual drones, self-driving cars, or autonomous infrastructure. An autonomic swarm framework can provide collective safety assurance even when problems emerge with an individual. This framework could build on research in zero-trust infrastructure, secure mesh networks, drone control systems, and collective intelligence. It could also inform better tools for testing, ensure transparent algorithms, and improve swarm alignment.

Sterritt has proposed combining the self-management aspects of autonomic computing and pre-programmed death as a safety mechanism through “apoptotic” computing to protect interconnected autonomous systems.[16] (“Apoptosis”—from the Greek for “falling off”—is the biological process of programmed cell death that purges cells that have completed their function or become diseased). More consideration must be given to integrating solutions across a range of possibilities to build trustworthy and assured autonomous systems. He concluded, “Without the development of such an approach, we will simply rediscover the risks of feature interaction at a higher level and in a way that is so dynamic as to be resistant to debugging and testing.”

Here are some important elements that could allow us to engineer trustworthy autonomous swarms for the next generation of drones, autonomous fleets, and robots:

Zero-Trust

Zero-trust security frameworks have evolved to protect the integrity and trustworthiness of decentralized systems. Cross-disciplinary TII research has explored how these might be extended to autonomous embedded systems. New physical and virtual testbeds are helping to identify issues and make improvements much earlier in the development of drone hardware, software implementations, and communications choices.

The next level of this research is exploring how to extend these principles to sensors and control systems. These investigations will guide development of a chain of trust relating to system malfunctions, faulty sensors, AI hallucinations, and drone attacks. Implementation of zero- trust security principles ensures that each drone in the swarm can be trusted to act safely and predictably.

Autonomic Networking

A drone swarm must ensure constant communication between individuals in the event of faulty equipment or cyber-attacks. This needs to support multiple communication frequencies and modalities to allow adaptation in the event of faulty radios, jamming, or poor network receptivity. TII is working on one approach for a secure mesh shield to ensure communication at the individual level, and which can extend the reach of communication when daughter drones fly outside the range of a mother drone.[17] [18] Furthermore, it provides failover to light communication when radio signals are jammed or damaged.

Scalability

It's also important to consider different ways of scaling drone swarms for different kinds of missions. At one extreme, a more decentralized approach might imbue each individual drone with an equal vote in overall control decisions. But this comes at the cost of configuring each device with the same high level of processing and communication as the others. At the other extreme, one centralized base station may coordinate the operations of all drones. This can cause problems, however, when any daughter drone flies outside the range of the base station, or when the mother drone is lost or damaged. A more pragmatic approach might balance a swarm of a few high-performance mother drones with a larger array of inexpensive daughter drones. If one mother fails, control could pass to adjacent mothers. This approach could improve scalability and resiliency while keeping costs down.

Additionally, for mission-critical applications, we need to develop highly available, secure, resilient, and safe ground station solutions—an effort that is also currently underway at TII.

Collective Intelligence

We also need to develop new consensus mechanisms for improving the trustworthiness and safety of drone swarms. When one drone perceives an anomaly, the others can corroborate the veracity of the data or identify a fault in the individual's perception. Similar principles could also be extended to the aberrant behavior of an individual. For example, if one daughter drone in a hierarchical organized swarm system behaves erratically, the mother drone could take control to keep it from crashing into a crowded area.

Testing and Validation

It's also important to develop comprehensive testing protocols and simulations to model the behavior of drone swarms. The TII is developing drone testbeds to characterize the emergent behavior of drone swarms in different environments, failure modes, and attack scenarios. The next step is to capture these behaviors in digital twins to help identify new problems and countermeasures to ensure safe operation in simulated environments.

Transparency

Innovations in AI and machine learning are helping to create surrogate models that perform thousands or millions of times better and more efficiently than symbolically coded models. These models show great promise in reducing the size and cost of control algorithms for individual drones. They can also fail in novel ways, leading to problems like hallucinated obstacles or mirage conditions that could imperil drone safety.

One big challenge is that some of the more efficient models come with millions or billions of parameters that do not directly correlate with human experience. Improvements in AI transparency and explain ability are required to correlate aberrant perceptions and behavior with specific algorithms or parameters to ensure their safety.

Alignment

It's also important to consider how these complex models interpret our goals to ensure alignment. These systems may pursue counterintuitive behavior that risks safety in pursuit of a high-level goal. The classic example is Nick Bostrom's paper clip maximizer thought-experiment, in which an AI system destroys the world while trying to create more paper clips.[19] More recently, a US Air Force thought experiment stirred global controversy when it hypothesized a drone might try to kill an operator who thwarted the machine's efforts to attain its assigned goal.[20] Large-scale digital twins that capture physical behavior and AI models could be essential in identifying and preventing these behaviors sooner in the development lifecycle.

Planting The Seeds For Efficiently Scaling Safer Swarms

The scientific and engineering communities have made tremendous progress in building safer autonomous drones, cars, robots, and devices of all kinds. The next step is to extend these concepts to swarms of collaborative autonomous things.

This new framework needs to consider an orchestration layer that automatically supports consensus building, resilient networks, and reliable sensor fusion across individuals and subgroups. The framework must include the entire drone stack, spanning chips, networks, and control systems.

We don't yet have many real-world examples to build on. Science fiction has presented two models to consider, though. At one extreme, lies Skynet control system of the Terminator series, a centralized controller guiding the movement of each individual. At the other extreme, there's Star Trek's Borg collective (before the Queen), a fully decentralized, interconnected mesh network of individuals working together. Neither scenario went well for humans; neither system aligned at all with human interests.

A more pragmatic approach might be a hybrid network spanning different kinds of systems working together and overseen by transparent control systems. A mix of high-performance parent controllers guiding inexpensive offspring could help these systems to scale. Innovations in AI alignment research will help ensure these larger swarms operate safely and as intended. Ultimately, we must build a safer model for humans, robots, and drones to work together at scale.

Our goal at SSRC is to invite researchers to collaborate, to drive development of secure, resilient, and safe autonomic systems that will transform the world into a safe place with new opportunities and economic growth, while protecting the environment and all the life on the planet.

Shreekant (Ticky) Thakkar and George Lawton were the key contributors. We like to thank Jean Pierre Giacalone and others who have inspired on this conversation.

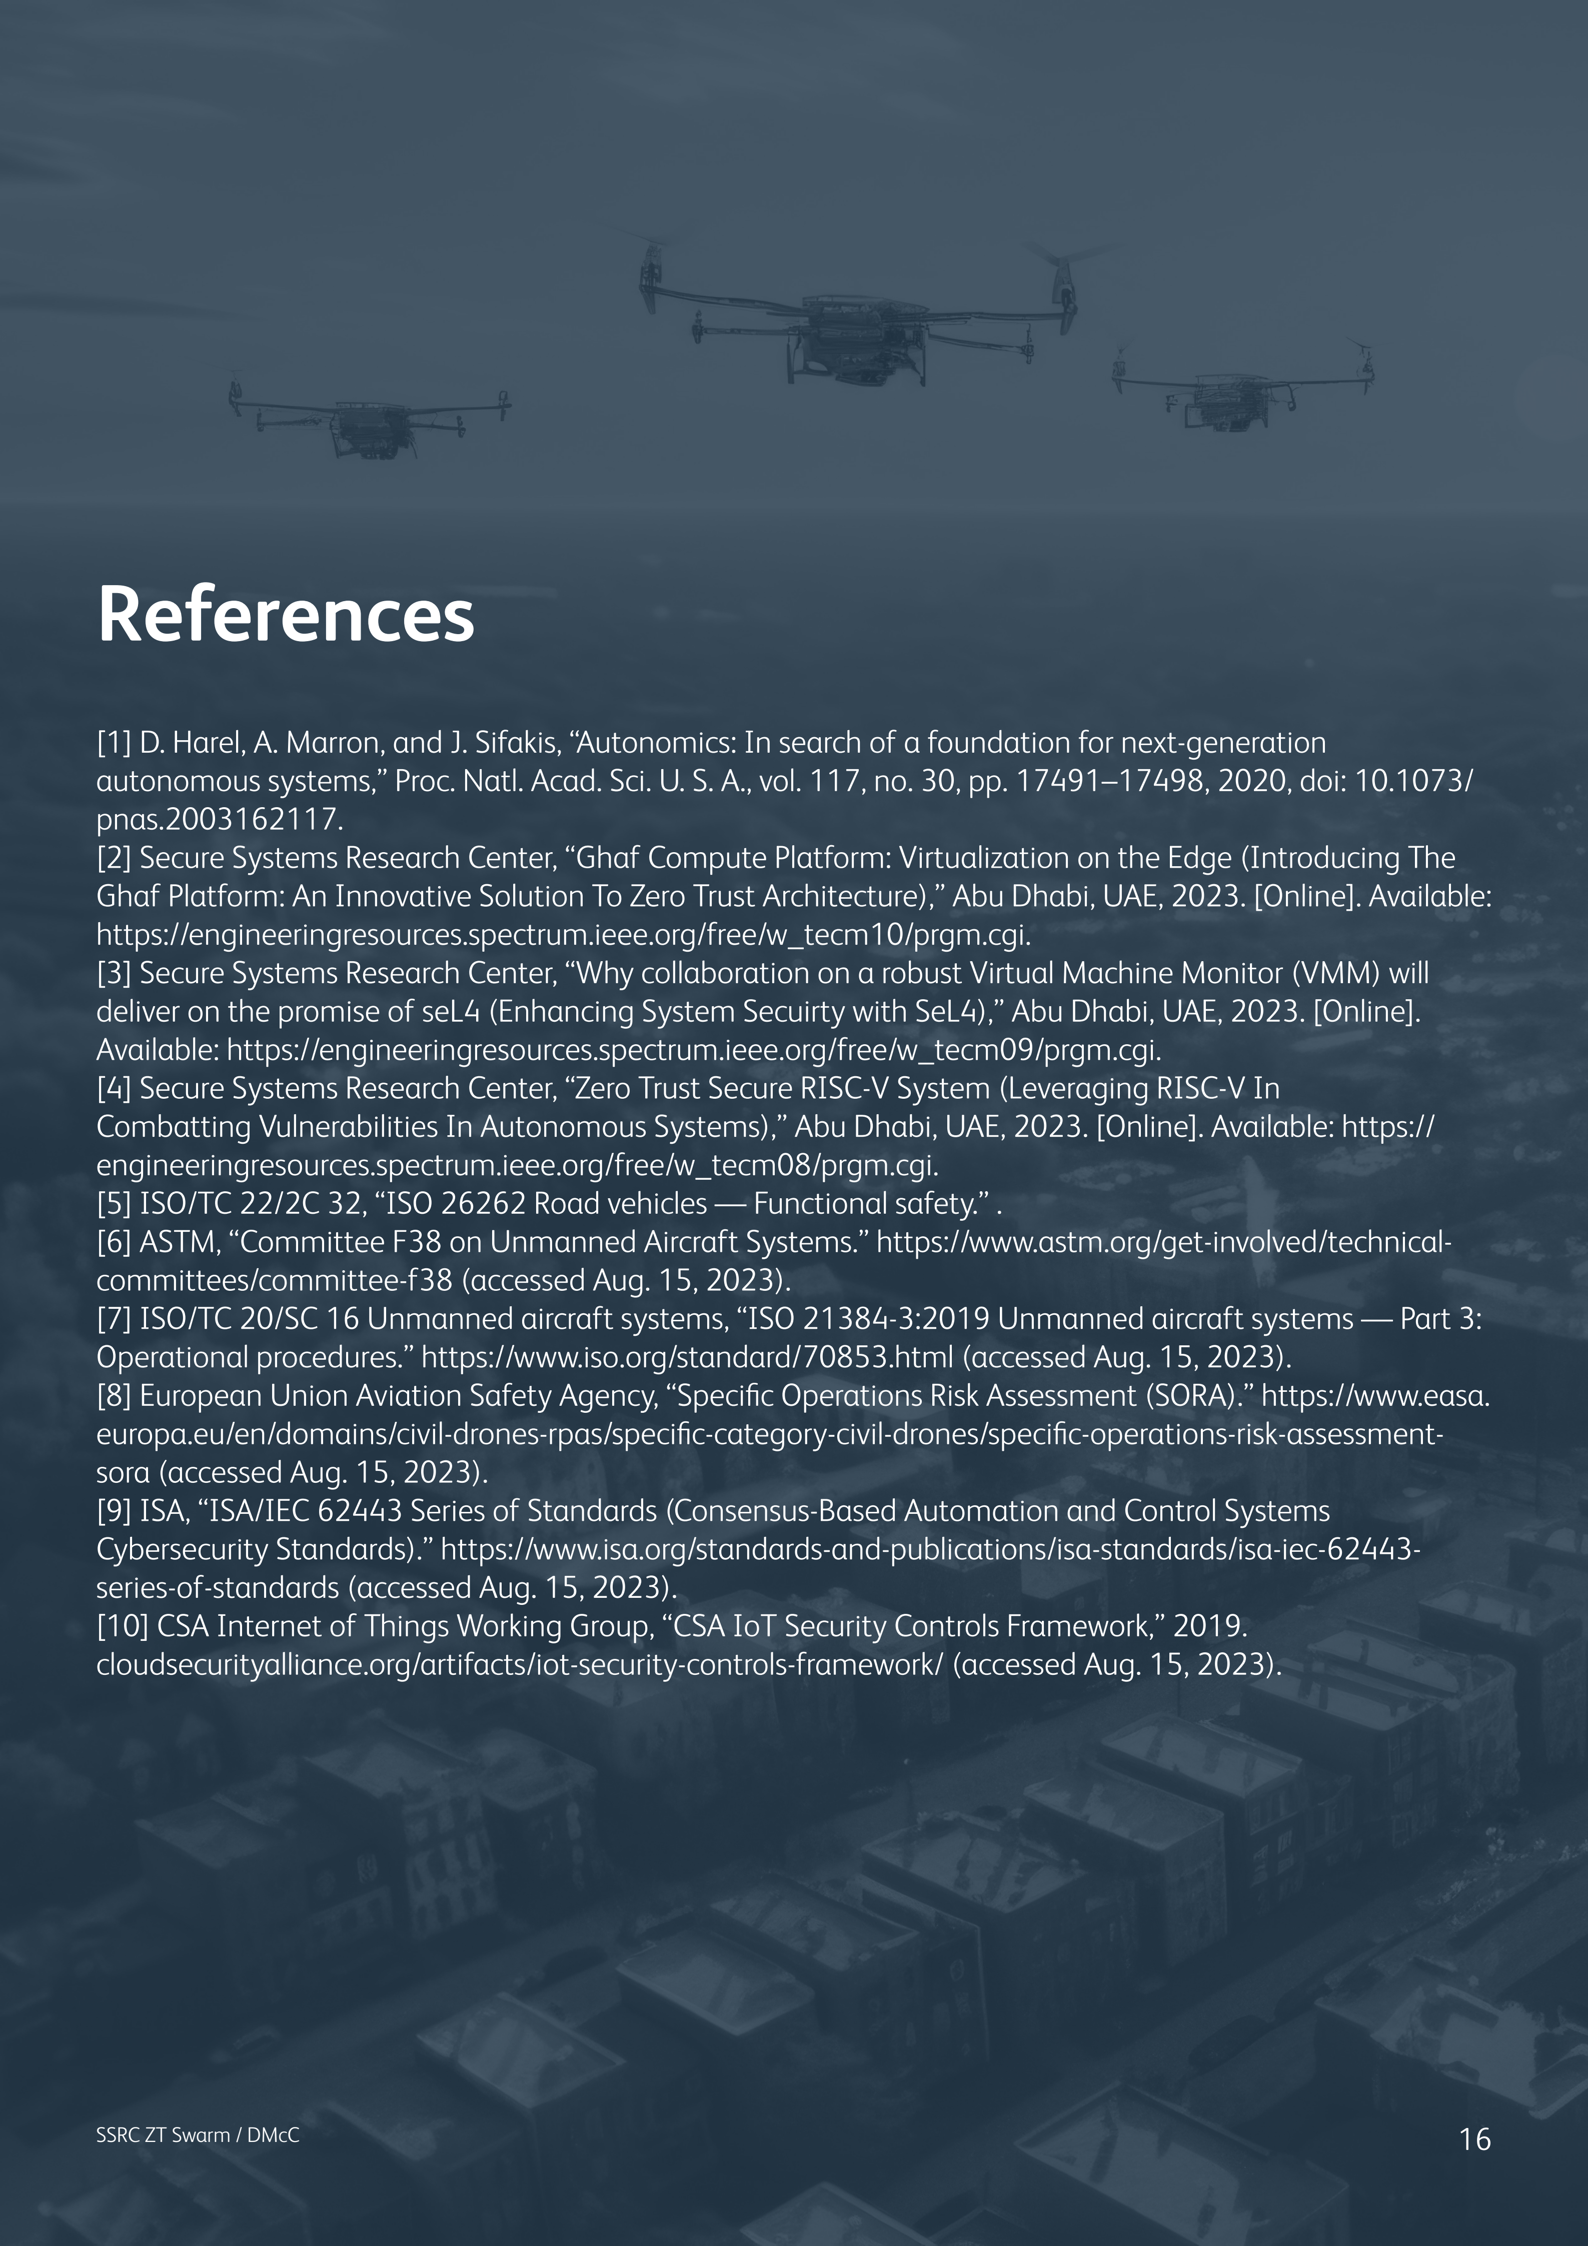

References

[1] D. Harel, A. Marron, and J. Sifakis, "Autonomics: In search of a foundation for next-generation autonomous systems," Proc. Natl. Acad. Sci. U. S. A., vol. 117, no. 30, pp. 17491–17498, 2020, doi: 10.1073/pnas.2003162117.
[2] Secure Systems Research Center, "Ghaf Compute Platform: Virtualization on the Edge (Introducing The Ghaf Platform: An Innovative Solution To Zero Trust Architecture)," Abu Dhabi, UAE, 2023. [Online]. Available: https://engineeringresources.spectrum.ieee.org/free/w_tecm10/prgm.cgi.
[3] Secure Systems Research Center, "Why collaboration on a robust Virtual Machine Monitor (VMM) will deliver on the promise of seL4 (Enhancing System Secuirty with SeL4)," Abu Dhabi, UAE, 2023. [Online]. Available: https://engineeringresources.spectrum.ieee.org/free/w_tecm09/prgm.cgi.
[4] Secure Systems Research Center, "Zero Trust Secure RISC-V System (Leveraging RISC-V In Combatting Vulnerabilities In Autonomous Systems)," Abu Dhabi, UAE, 2023. [Online]. Available: https://engineeringresources.spectrum.ieee.org/free/w_tecm08/prgm.cgi.
[5] ISO/TC 22/2C 32, "ISO 26262 Road vehicles — Functional safety." .
[6] ASTM, "Committee F38 on Unmanned Aircraft Systems." https://www.astm.org/get-involved/technical-committees/committee-f38 (accessed Aug. 15, 2023).
[7] ISO/TC 20/SC 16 Unmanned aircraft systems, "ISO 21384-3:2019 Unmanned aircraft systems — Part 3: Operational procedures." https://www.iso.org/standard/70853.html (accessed Aug. 15, 2023).
[8] European Union Aviation Safety Agency, "Specific Operations Risk Assessment (SORA)." https://www.easa.europa.eu/en/domains/civil-drones-rpas/specific-category-civil-drones/specific-operations-risk-assessment-sora (accessed Aug. 15, 2023).
[9] ISA, "ISA/IEC 62443 Series of Standards (Consensus-Based Automation and Control Systems Cybersecurity Standards)." https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards (accessed Aug. 15, 2023).
[10] CSA Internet of Things Working Group, "CSA IoT Security Controls Framework," 2019. cloudsecurityalliance.org/artifacts/iot-security-controls-framework/ (accessed Aug. 15, 2023).

[11] E. U. A. for C. (ENISA), “Baseline Security Recommendations for IoT,” 2017. https://www.enisa.europa.eu/publications/baseline-security-recommendations-for-iot (accessed Aug. 15, 2023).
[12] US Federal Aviation Administration, “Unmanned Aircraft System Traffic Management (UTM),” 2023. https://www.faa.gov/uas/research_development/traffic_management (accessed Aug. 15, 2023).
[13] D. B. Abeywickrama et al., “AERoS: Assurance of Emergent Behaviour in Autonomous Robotic Swarms,” Arxiv, 2023, [Online]. Available: http://arxiv.org/abs/2302.10292.
[14] C. Macrae, “Learning from the Failure of Autonomous and Intelligent Systems: Accidents, Safety, and Sociotechnical Sources of Risk,” Risk Anal., vol. 42, no. 9, pp. 1999–2025, 2022, doi: 10.1111/risa.13850.
[15] S. Mann, “Tesla driver blames self-driving mode for eight-vehicle pile-up on San Francisco’s Bay Bridge,” Daily Mail Online, Dec. 23, 2023.
[16] R. Sterritt, “Achieving Trustworthy Autonomous Systems through Autonomic and Apoptotic Computing: Keynote at ICAS 2020,” pp. 53–58, 2020, Accessed: Aug. 15, 2023. [Online]. Available: https://pure.ulster.ac.uk/ws/portalfiles/portal/86212464/20054.pdf.
[17] S. (Ticky) Thakkar, “A Zero Trust approach to autonomous systems of systems security,” Abu Dhabi, UAE, 2023. [Online]. Available: https://engineeringresources.spectrum.ieee.org/free/w_defa3180/prgm.cgi.
[18] Secure Systems Research Center, “Secure Comms Shield: A Zero-Trust Network Solution for Distributed Communication (Zero Trust Architecture in Combatting Wireless Network Attacks),” Abu Dhabi, UAE, 2023. [Online]. Available: https://engineeringresources.spectrum.ieee.org/free/w_tecm07/prgm.cgi.
[19] “Instrumental convergence,” Wikipedia. https://en.wikipedia.org/wiki/Instrumental_convergence (accessed Aug. 15, 2023).
[20] S. Sharma, “Air Force official ‘misspoke’ in tale of AI drone killing human operator in US test mission,” Independent, Jun. 03, 2023.
[21] https://www.fox7austin.com/news/autonomous-vehicles-chaotic-scene-west-campus-austin-texas

Technology Innovation Institute LLC
P.O. Box 9639, Abu Dhabi, UAE

tii.ae