Security-per-Watt: Making Zero Trust Work When You Have Three Watts and Intermittent Comms

Jonathan Shelby
Department of Computer Science, University of Oxford

Zero Trust Architecture (ZTA) has become the dominant security paradigm for enterprise networks. The core premise — never trust, always verify — demands continuous authentication, fine-grained access control, and pervasive monitoring. These requirements are tractable when backed by cloud-scale infrastructure. But what happens when the system you need to secure runs on a single-digit watt power budget, communicates through intermittent satellite links, and cannot be physically accessed for maintenance? This is the reality for CubeSat constellations and a growing class of constrained edge platforms, and current ZTA guidance has nothing useful to say about it.

This talk presents Security-per-Watt (SpW), a heuristic framework developed to reason about achievable security assurance as a function of available power, compute, and communications budgets. Rather than treating security requirements as fixed and asking whether a constrained platform can meet them, SpW inverts the question: given a concrete resource envelope, what is the maximum-security assurance a system can deliver, and where should it spend its budget for the greatest defensive return?

The framework was developed and evaluated in the context of CubeSat platforms, where the constraints are stark. A typical 3U CubeSat operates on roughly 5-10W of generated power, runs a radiation-hardened processor several generations behind commodity hardware, and may have connectivity windows of minutes per orbital pass. We analysed how standard Zero Trust components — identity verification, policy evaluation, telemetry collection, encrypted communications — map onto this resource envelope, and where the binding constraints bite. The results show that meaningful security guarantees are achievable even under severe resource limitations, but only if the system is deliberate about which verification steps it prioritises. Blanket application of enterprise ZTA patterns fails; selective application guided by SpW analysis does not.

A key finding is that the trade-off between security coverage and resource consumption is not linear. Some verification steps are cheap and high-value; others consume disproportionate resources for marginal security benefit. This asymmetry creates an opportunity for principled optimisation that ad hoc engineering misses. We propose a Distributed Security Paradigm (DSP) that pushes security decision-making to the network edge, reducing dependence on centralised policy infrastructure that constrained platforms cannot reliably reach.

This work was completed as part of an NCSC-certified MSc at Oxford and forms the foundation for ongoing DPhil research. The natural next step is to generalise SpW into a broader Security-per-Resource framework and to place the optimisation problem on formal foundations — potentially using game-theoretic models to capture the adversarial dynamics of resource-constrained security allocation. I am at an early stage with this generalisation and would welcome the community's perspective on

two questions: whether the constrained-resource framing resonates beyond satellites and IoT, and what formal methods the systems community considers most promising for reasoning about security-resource trade-offs in practice.