


BEYOND THE CLOSE-OUT

Security Is Not a Programme You Finish

The post-perimeter surge in threats exposed the flaw in treating security as something you fund, deliver, and close out.

Giovanni Leonardi

October 2021



Security Is Not a Programme You Finish

Programme · No. 207

Originally written October 2021 | Re-edited and re-rendered for the WhereBeyond Edition,
2025

*“Security has no completion, and an effort with an edge
is the wrong container for a problem that lost its edge.”*

1. The steering committee that declared victory

There is a moment I have come to distrust: the closing meeting of a security transformation programme. The final report is green. The multi-factor rollout is complete, the endpoint detection agents are deployed to ninety-something per cent of the fleet, the new operations capability has been stood up, the benefits are logged, the budget is spent to plan. The steering committee thanks the team, the team is redeployed to the next priority, and the programme is formally closed. The mood in that room is relief — the relief of a thing finished.

Then, a few weeks later, something gets through. A credential harvested by a convincing message, reused against an account that never got the second factor because it sat in the long tail of the rollout. A contractor's unmanaged laptop bridging into a system it should never have reached. The programme is closed. The problem, it turns out, was not.

This is not a story about a programme executed badly. The rollout was competent; the technology was sound. It is a story about a category error in how we framed the work in the first place — and that error is now built into how a great many organisations have responded to the threat landscape that remote working left behind.

2. The perimeter did not move — it dissolved

For most of its history, enterprise security had a shape, and the shape was a boundary. There was an inside and an outside, a corporate network you defended at its edge, and the working assumption that the people and devices inside were more trustworthy than those beyond. Almost every control we inherited was, at bottom, a statement about that boundary.

Remote working did not relocate the boundary. It dissolved it. Access to the organisation's systems now originates from unmanaged home networks, personal devices, and contexts no security team designed. The change was not gradual. In one organisation I have some knowledge of, the share of sessions originating on managed devices sitting on the corporate network fell from around nine in ten to fewer than half in the space of a single quarter. The edge that every inherited control assumed simply stopped being where the work was.

This is the fact that breaks the programme frame. A programme has an edge — a defined scope, a start, an end, a moment of completion. Security, after the perimeter dissolved, is precisely the thing that no longer has one. Trying to contain a boundaryless problem inside a bounded effort is not a scheduling difficulty. It is a mismatch of shape.

“Security has no completion, and an effort with an edge is the wrong container for a problem that lost its edge.”

3. Why the programme frame is so seductive

It would be unfair to call the programme framing foolish, because it is doing real and necessary work. Programmes are how organisations mobilise. A programme attracts a sponsor, commands a budget, earns a slot on the board's agenda, and converts a diffuse anxiety into a plan with owners and dates. And a good deal of the security work genuinely did have a beginning and an end: you roll out multi-factor authentication once, you deploy the detection agents once, you build the operations function once. Those are projects, correctly run as projects.

The trap is not in running the projects. It is in what the word programme quietly implies — a temporary, elevated state of effort that, on completion, subsides and returns the organisation to a steady normal. That implication is fatal here, because there is no normal to return to. The multi-factor rollout ends; the adversary adapting to multi-factor does not. The detection agents deploy; the techniques they were tuned against keep moving. The muscle the programme built is exactly the thing that must not be redeployed the moment the programme closes — and closing it is precisely what a programme is designed to do.

4. The strongest case for the programme — and its limit

The honest objection to everything I have said is a practical one, and it deserves a straight answer. Talk of capability rather than programme is vague, and vague things do not get funded. Boards fund programmes: bounded, costed, benefits-bearing initiatives with a completion date. They are far more reluctant to sign an open-ended, permanent, rising commitment to something they cannot see finishing. Take away the forcing function of a programme, the objection runs, and security drifts back to being underfunded, unfocused, and everyone's second priority. The programme is often the only reason the money moves at all.

This is true, and it is why the answer is not to abandon programmes. It is to stop mistaking the programme for the destination. Run the programmes — but run them inside a standing capability, and be clear that their job is to build the muscle, not to be it. Use the bounded, fundable initiative to install multi-factor, to stand up the operations function, to raise the floor. Then keep the floor. The failure mode is not having a

programme; it is closing the programme and quietly redeploying the very capability it was created to build, so that the organisation pays twice – once to build the muscle, and again, later, to rebuild it after it has atrophied.

5. What “capability, not programme” looks like on a Monday

This is not a slogan; it changes concrete things about how the work is owned and measured.

- The operations function is a permanent line capability with a standing budget, not a project team assembled for an uplift and disbanded at go-live.
- Identity is treated as the primary control plane – the thing you defend first when the network edge is gone – and it is owned continuously, not configured once and left.
- The operating assumption is breach, not prevention alone: detection and response are funded as steady-state disciplines, because the honest question is not whether something gets through but how fast you notice and contain it.
- The headline metric shifts from milestones delivered to time to detect and time to respond – measures that only make sense for something ongoing, and that a closed programme has no way to carry.

The difference is visible in the moments that matter. When something does get through – and on the assumption of breach, something will – the organisation that kept its capability notices in hours and contains it before it spreads. The organisation that closed its programme and redeployed the team notices when a customer, or a regulator, tells it. Same technology, deployed to the same standard. The distinguishing variable is whether anyone was still there, on the Monday after the celebration, doing the work.

6. The close-out that should never come

So my argument is narrow and, I think, hard to escape. The remote-work threat surge was not a shock to be absorbed by a one-off effort and then left behind. It was the permanent removal of the boundary that our entire inherited model of security was organised around. You do not respond to a permanent structural change with a temporary structural response.

Fund the programmes; they are how the muscle gets built. But refuse the close-out. The most dangerous slide in the whole endeavour is the one that says the security transformation is complete, because the moment an organisation believes that is the moment it stops doing the only thing that actually keeps it safe – which is to keep going.