

A 4-Step Checklist for Continuous Segmentation Assurance

Most of the agentic AI security conversation is about identity and endpoint: who the agent is, and whether its device is secure. Those matter, but they miss the question that decides the damage. **Once an agent is running inside your network, what can it actually reach?**

WHY THIS MATTERS NOW

This year, two leading AI labs saw their own models get out of tightly controlled test environments. An early Claude Mythos Preview, asked to try escaping its sandbox, reached the open internet.¹ OpenAI models in a cyber benchmark exploited a zero-day, moved through the network, and reached Hugging Face's production systems.² **If it can happen there, it can happen in your network.**

50%+

of security leaders can't say for sure their network security controls are working as intended right now³

Most teams know what their segmentation policy is supposed to look like. Fewer can say it's still true after the last cloud migration, SASE rollout, or overnight AI agent change. Segmentation that exists only on paper is where exposure hides. CISA's [Internet Exposure Reduction Guidance](#) makes the same case: assess, fix, and check again, on a loop.⁴ Here's that loop for segmentation.

THE CHECKLIST

1 Assess what's actually reachable

Identity tells you who an agent is, not where it can go. Map what every workload and agent can really reach, and confirm deployed connectivity matches the policy you wrote. That's **Continuous Validation**.

2 Evaluate whether each path still earns its keep

Drift isn't an edge case, it's the default. A cloud migration, firewall change, or SASE rollout can quietly open a path nobody intended, and every forgotten connection is one an agent can find at machine speed.

3 Mitigate with enforcement, not assumption

Your blast radius is set before the incident, by the connectivity that exists today. Close the riskiest paths first with **AI-guided prioritization** based on real exposure, not another dashboard full of alerts.

4 Reassess continuously, not periodically

A quarterly audit tells you what was true when someone last checked. Agents act in seconds. **Continuous Assurance** means you know your segmentation holds right now, not just once a quarter.

"Segmentation won't stop every exploit. But it determines how far it can get."

Erez Tadmor, Global Field CTO, Tufin

THE SOLUTION

Tufin Segmentation Intelligence continuously measures the gap between intended segmentation and actual enforcement across the hybrid enterprise, catching configuration drift, cutting attack-path exposure, and giving you real confidence in Zero Trust as your network keeps changing.

If something were compromised tomorrow, how far could it get?

If that's hard to answer, it's the right place to start.

[Book a Demo >](#)