PrimusCredence

Provable Security for Applications

A new field beside your AppSec tools — what cannot happen in your system.

1. **Access control** — who may do what
2. **Network isolation** — what can reach where
3. **Data residency** — where data can go

primuscredence.com

UAE

Dr. Raghavendra Ramesh

Founder & CEO, PrimusCredence

22+ YEARS OF EXPERIENCE IN PROVABLE SECURITY

PhD, IISc Bangalore
Model checking for information-flow security

Oracle Labs · 2014–2019
Java vulnerability detection

ConsenSys R&D · 2019–2021
Cross-chain protocols, consensus verification

Supra · VP of R&D · 2021–ongoing
BFT design & FV, bridges, DeFi & Crypto

TECHNIQUES

Static analysis, theorem proving, model checking

TOOLS

Dafny, TLA+, IVy, Lean, Isabelle

Where the Risk Actually Lives

The cloud base is proven. The app above it — and the seams between — are not.

Third-party plugins & extensions 1
Run inside the app — and inherit its privilege

APPLICATION 1

Its own roles, API tokens and public users — the authorisation logic and configuration in its database and admin panel

OWASP #1 Broken Access Control · #2 Misconfiguration — in 100% of apps

THE APP ↔ CLOUD & DATA SEAM

1 Access via APIs 2 Network paths & zones 3 Data paths & copies

CLOUD & INFRASTRUCTURE

VMs · containers · databases · queues · object storage · IAM · network

Access control & network — formally verified by the providers

UNPROVEN TODAY

Where the breaches are

The app's own access control, its plugins, and the access, data and network seams beneath it.

→ Services 1 · 2 · 3 prove it.

PROVEN BASE

Cloud IAM & network — already formally verified.

OWASP 2025 — Access Control & Misconfiguration

The two categories at the top are found in 100% of apps tested.

A01

Broken Access Control

#1 — unchanged

40 CWEs · 1.84M occurrences · 32,654 CVEs

SSRF now folded in — easy to exploit, widespread.

A02

Security Misconfiguration

#2 — up from #5

16 CWEs · 719,084 occurrences

Again, 100% of tested apps carried a form of it.

BROKEN ACCESS CONTROL IN THE WILD

Capital One	106M	WAF SSRF → over-privileged role; a tipster found it, ~4 months in
Optus	9.5M	access-control bug on a dormant, internet-facing API — latent ~4 years
T-Mobile	37M	API data retrieval without authorization — ~6 weeks to notice
Dell	49M	auto-approved fake partner, then service-tag enumeration
Trello	15M	unauthenticated email lookup through rotating proxies
Twilio Authy	33M	unauthenticated endpoint enumerated for phone numbers

Each system did exactly what its access rules allowed — the rules were wrong.

Access Policy Lives in Many Places

Each place has its own owner and its own review. Nobody checks the join.

ONE REQUEST ---------->

Identity provider

Groups and claims — who is a doctor, who audits

IAM team

IN AN ENGAGEMENT

API gateway / router

Which routes are open, and to whom

Platform team

IN AN ENGAGEMENT

Admin panel

The role grid — “staff may read documents”

App admin

IN AN ENGAGEMENT

Application code

Middleware, controllers, plugin resolvers

Developers

IN AN ENGAGEMENT

Database

Grants, row-level rules, tenant fields

DBA

IN AN ENGAGEMENT

Cloud IAM role

What the app itself may touch — buckets, keys, queues

Cloud team

IN AN ENGAGEMENT

Often the code is the only written policy: MACE (CCS 2014) found access-control flaws in 5 of 7 real web apps.

An engagement learns every place your policy lives — and checks the join.

Segmentation Is Asserted, Not Proven

Once inside, lateral movement turns an incident into a breach.

29 min
average eCrime breakout in 2025: initial access to lateral movement

27 sec
the fastest breakout ever observed

82%
of detections in 2025 were malware-free

Source: CrowdStrike 2026 Global Threat Report (2025 data)

WHY PROBING CAN'T ANSWER IT

- Tests only the paths you thought of
- Can't tell blocked from about-to-change
- East-west is claimed more than shown

REACHABILITY IS A CONFIG PROPERTY

- Routes, SGs, NACLs, firewall rules
- Decided statically — no scanning window
- How hyperscalers prove their own

On OT networks active probing is often off-limits — the IT/OT boundary must be proven.

Residency Is Reviewed, Not Proven

Every change is reviewed alone — nobody re-checks where the data can now go.

FOUR QUIET WAYS OUT

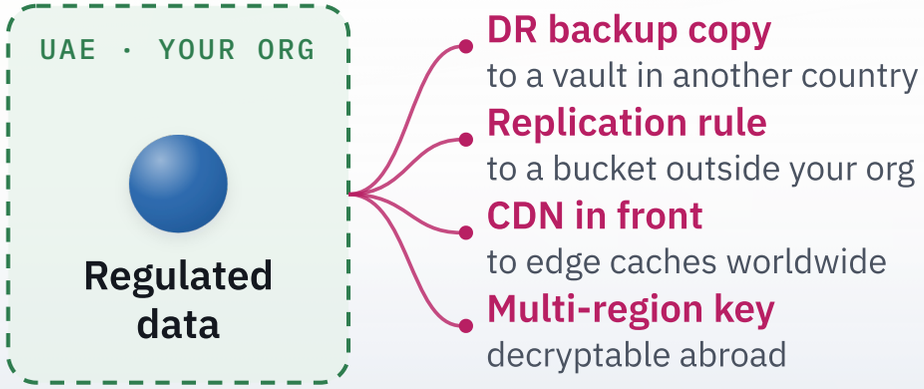

WHY THE TOOLS STAY GREEN

- Region locks govern where, not whose
- DSPM shows where data is, not where it can go
- CSPM checks one resource at a time — and rewards cross-region backups

€1.2B

Meta, 2023: EU data sent to the US (Irish DPC)

€530M

TikTok, 2025: EEA data accessed from China (Irish DPC)

€290M

Uber, 2024: EU drivers' data kept in the US (Dutch DPA)

6 flaws

AWS, 2024: guessable bucket names leaked data to attackers

A region lock answers where. Sovereignty also asks whose.

AI Industrialises the Attacker

Point-in-time testing cannot keep pace with an enumerating attack.

54M
lines of code analysed by autonomous AI systems (DARPA AIxCC 2025)

$3.52
compute cost per LLM-agent exploit attempt on a real CVE (Fang et al. 2024)

87%
of 15 real CVEs exploited from the public advisory alone (Fang et al. 2024)

TRADITIONAL PEN TEST	AI-DRIVEN ADVERSARY
10^2–10^3 paths, human-selected	10^5–10^7 paths, machine-enumerated
2–4 week window, then closed	Continuous, no window
High marginal cost per path	Near-zero marginal cost
“We found no way in”	Approaching exhaustive

Path counts are orders of magnitude, for illustration.

Proof is the one capability that answers exhaustive search in kind.

What Today's Security Stack Leaves Open

Each tool answers its own question well. One question stays open.

SAST	Knows code patterns, not intent. A check present but wrong still passes.
DAST	Tests the endpoints and personas it is given — not every user × object.
IAST / RASP	Same-code-path BOLA is invisible — was the call even permitted?
DSPM	Finds where regulated data is — never where it can go.
CNAPP / CSPM	Stops at the cloud IAM role; never reads the app's own roles.
WAF / API sec	Judges traffic, not grants: whether it should be permitted is upstream.
Pen test	Shows a way in, or that none was found — not that none exists.
SOC · SIEM	Sees what happens — not what the rules would allow next.

They all share one gap: no statement of what is allowed, decided over every request.
Keep every one of them — and add the new field that closes it: provable AppSec.

Proof — Now Feasible

Formal verification used to take expert teams years. AI brings it into range.

THEN

- Expert teams, bespoke specs, years
- Confined to avionics, silicon, protocols
- “End users won’t write a spec”

NOW — AI-ASSISTED

- AI assists modelling & proof search
- Expert owns & ratifies the spec
- AWS: ~1B reasoning queries a day

What we add A preventive control with evidence behind it, beside your testing and SOC.
We add to your stack; we don’t replace it. Proof alone isn’t enough.

WHAT STAYS WITH YOU — ALL OF IT NEEDED

Configuration drift, shadow assets and anything outside the model

Insider misuse, and valid credentials used exactly as intended

Detection, response and the rest of the programme you already run

Automated Reasoning

Proof establishes that bad things cannot happen.

1 — Learn
We learn the app's logic and turn it into a logical model

2 — Define
Its desired security properties become formulas

3 — Solve
Automated solvers decide every case

4 — Report
A proof — or the attack, a fix and a re-proof

AI ASSISTS EVERY STEP · THE EXPERT OWNS THE SPEC

PROOF

No violating state exists — over the entire space, under a stated model.

Not “none was found” — a strong compliance statement.

COUNTER-EXAMPLE

The exact request, path or chain that breaks it — reproducible on your systems.

A property fails? We demonstrate the attack existing tools miss, help developers fix it, and re-prove.

A proof holds only for the property and model stated · Industrial practice: AWS · Microsoft · Meta · Airbus.

Services

One method — three services: access control, network isolation, data residency.

1 ACCESS CONTROL

App entitlement proofs

Every app role, token, plugin route and public user — decided over all requests.

Who may do what

2 NETWORK ISOLATION

Network segmentation

Routes, security groups, firewall and Kubernetes rules — every path between zones.

What traffic can reach where

3 DATA RESIDENCY

Residency proofs

Backups, replicas, keys, caches — not where data is, but where it can go.

Where regulated data can go

EVERY SERVICE ENDS in a proof — or the exact request, path or chain that breaks it.

1. Application Access-Policy & Entitlement Verification

> *Can any app role, API token, plugin route or public user reach data it must not — through the app or its cloud role?*

THE PROBLEM

Policy is scattered: admin grid, code, plugins, identity groups, cloud role. Often the code is the only spec. Each part passes review; nobody checks the join.

An engagement learns every place your policy lives — and checks the join.

WE READ	Roles, permissions, routes, cloud role
ACCESS	Read-only, nothing installed
BUYERS	Platform & AppSec teams · ISVs

IN THE DEMO

tools run first; none decides these

VIOLATED **One setting overrides correct code**
One admin setting exposes another tenant's data.

VIOLATED **Two allowed steps, one breach**
Each step is allowed; together they cross tenants.

VIOLATED **Side doors**
A plugin route and the cloud role bypass the app.

PROVED **after each fix — every user, every request**

OUTPUT A proof that no role, token, plugin route or public user reaches what it must not — or the exact request that breaks it.

SERVICE 1 — THE ENGAGEMENT

What We Ask of You

A consulting engagement — we work alongside your developers to improve security.

1

Full read-only access

To the code and the cloud instances — every component: the app, its plugins, identity and cloud roles.

The join is where the gaps hide

2

App developers

A named developer point of contact, and periodic meetings with the team through the engagement.

We fix alongside them, not at them

3

AI, on your account

Our AI tools run as an authorised user of your organisation, on your own AI account.

Your data never passes through ours

Your code, your cloud, your AI account — the work stays inside your organisation.

SERVICE 2 — WHAT TRAFFIC CAN REACH WHERE

2. Network Reachability & Segmentation Proofs

> *Does any path exist from an untrusted zone to this one — and does east-west segmentation actually hold?*

THE PROBLEM

Probing tests the paths you thought of. Reachability is a property of the configuration — routes, firewall rules, NAT — across every segment at once.

No scanning window, no production traffic — safe on OT, and re-run on every change.

WE READ	Device and cloud network configs
ACCESS	Read-only, no packet sent
BUYERS	Finance · OT & CNI · sovereign cloud

IN THE DEMO

tools run first; none decides these

VIOLATED **One policy, every zone pair**
Cloud into OT, hop by hop — each rule, each owner.

WOULD VIOLATE **A change, before approval**
A pending firewall rule would open paths into OT.

VIOLATED **A one-character slip**
Any workload reaches another tenant's database.

PROVED **every row, after the fix — no packet sent**

OUTPUT A proof that no path exists — or the path, hop by hop, with the rule that permits it and the owner of that rule.

SERVICE 3 — WHERE REGULATED DATA CAN GO

3. Data Residency & Sovereignty Proofs

> *Can regulated data leave the country or the organisation — or could anyone make it leave?*

THE PROBLEM

Residency is asserted by configuration review. Backups, replicas, caches and keys are each reviewed alone. DSPM shows where data is, not where it can go.

Not where your data is — where it can go.
Today, and whatever anyone could set up.

WE READ	Cloud, storage, key and CDN policies
ACCESS	Read-only, nothing installed
BUYERS	Sovereign cloud · government · banks

IN THE DEMO

tools run first; none decides these

AT RISK **Nothing has left — yet**
An operations role could move the data outside the organisation.

VIOLATED **Copies abroad**
A backup and edge caches sit outside the country.

OUTSIDE MODEL **Application egress**
Shown, not hidden: the app can reach the internet.

PROVED **for storage, keys and edge after the fix**

OUTPUT A proof that no channel moves the data outside the approved jurisdiction or organisation — or the exact chain that would move it.

Our Approach

We partner with cybersecurity solutions providers, adding provable assurance.

HOW WE PARTNER — THROUGH THE PROVIDERS

Provable control effectiveness inside their bids.

They keep the client relationship and the contract.

We don't certify — a channel, not a rival.

WHOM WE REACH THROUGH THEM — GCC GOVERNMENT & PRIVATE ENTERPRISE

Government & sovereign	Federal & emirate entities · Digital-government programmes
Aviation & CNI	Air navigation · Airports · Utilities
Energy, oil & gas	National energy companies · OT operators
Banking & finance	Regional banks · Sovereign funds · FinTechs
Telecom & cloud	Telecom operators · Sovereign-cloud providers

Why the Provider's Clients Will Ask

Four things happened at once in the GCC — and their clients feel all four.

Mandatory replaced voluntary

UAE PDPL is federal law; SAMA CSF and NCA ECC bind banks, government and CNI.

Assurance shifted to evidence

Regulators increasingly ask for demonstrable control effectiveness, not attestation.

Third-party risk is now in scope

Rules increasingly require evidence that integrations and components are controlled.

Sovereignty must be demonstrable

Localisation asks whether regulated data can leave — our residency proofs settle it.

WHERE A FINDING ALSO LANDS

DESC ISR	1 2 3	NESA / UAE IAS	2 3	UAE PDPL	2 3
SAMA CSF	1 2 3	NCA ECC	1 2 3	PCI DSS 4.0	1 2
ISO 27001	1 2 3	CIS Controls	1 2 3	NIST CSF 2.0	1 2 3

Regulation routes the finding to a second reader. It is not the reason to commission it.

What a Provider Can Resell

We do not certify — so an audit firm or an MSSP is a channel, not a rival.

Detection · response · posture reporting **+** **Provable control effectiveness**

THE COMMERCIAL MOVE

- Proof inside bids you already win
- Recurring: re-proven on every change
- You keep the client and the contract

WHY IT DEPLOYS EASILY

- Read-only, nothing installed on any host
- No scanning window, no production traffic
- Verdicts can land in your SIEM and tickets
- AI runs in your tenancy, data in-region

We produce the technical evidence. Your assessor keeps the opinion.

One Result, Three Registers

One analysis, read by engineering, risk and audit — each in its own terms.

First line

Engineering and security operations

The property, its model and assumptions, and either the proof or a reproducible attack path.

Second line

Risk

Likelihood moved off a qualitative guess, with residual risk and the model's limits stated.

Third line

Internal audit and compliance

A clause-mapped control assertion: objective, status, owner and re-verification date.

✓ Design → operating effectiveness · attaches to your GRC records · assessors probe less

The first register is the product. The second and third make it worth more than it cost.

Take Away

Bringing mathematical rigour to security —
proof alongside testing, from the app to the cloud.

The shift

Attack path discovery is automated, parallel and cheap. Detection shows what is happening; proof shows what cannot happen.

The addition

We do not replace testing, monitoring or the SOC. We make one class of controls machine checkable.

The honesty

A proof is relative to a model and a specification — both stated. We produce the evidence; we do not certify.

START HERE — MICRO ASSESSMENT OF SERVICE 1

One narrow, high-risk slice. Read-only access, nothing installed, weeks not months — and a counter-example on your own systems, or a proof.

Thank You

primuscredence.com

Dr. Raghavendra Ramesh
Founder & CEO, PrimusCredence
PhD, IISc Bangalore

raghavendra@primuscredence.com
+971 58 189 2803
UAE

Provable Security for Applications