



---

# Mythos – Myth or Reality!

What can you do right now, to prepare your organization for when Mythos becomes available to everyone.



**Anirban Mukherji**

anirban@miniOrange.com

**Founder & CEO, miniOrange**

# Overview & offerings ..Since Oct 2012

**30K+**  
CUSTOMERS

**800+**  
EMPLOYEES

**Pune**  
HEADQUARTERED



## Identity Security

Identity & Access Management (IAM)

Privileged Access Management (PAM)

Identity Governance & Administration (IGA)

Identity Threat Detection & Response (ITDR)



## Endpoint Management

Data Loss Prevention (DLP)

Mobile Device Management (MDM)

Cloud Access Security Broker (CASB)



## Data Privacy

Consent & Cookie Management

DPIA

RoPA

DSAR

Data Discovery



## CyberSecurity As A Service

SAST

DAST

Red Teaming

VAPT

Third-Party Security Assessment



## AI Security

Agentic AI

Non Human Identity

Governance & Observability

Policy Enforcement

SPRING 2025

**Momentum Leader**

SPRING 2025

**High Performer**

SPRING 2025

**Best Meets Requirements**

SPRING 2025

**Easiest To Use**



# Mythos - Myth or Reality

## MYTH - the sky is falling

✗ MYTH

Mythos will instantly hack every system - there's nothing you can do to stop it

✗ MYTH

Mythos is locked away safely - restricted access means we're all protected for now

✗ MYTH

AI creates brand new vulnerabilities that never existed before

✗ MYTH

AI attacks are unstoppable because they operate at superhuman speed

✗ MYTH

Traditional security tools are obsolete in the age of AI

## REALITY - basics still win

✓ REALITY

Mythos targets weak posture - MFA, least privilege, and patching stop most attacks cold

✓ REALITY

Similar capabilities are spreading to other labs and open-weight models - restricted access is a window, not a wall

✓ REALITY

AI exposes existing flaws faster - your unpatched CVEs are the real threat

✓ REALITY

Speed amplifies both sides - automated detection, rate limiting, and anomaly monitoring can shut down attacks just as fast

✓ REALITY

Your existing stack still matters - EDR, WAF, IAM, and SIEM become even more effective when tuned and integrated properly












**Good news:** You're not helpless. Smart basics. Right now. **Make all the difference.**



# Build-time hygiene (SBOM, patching dependencies)

This is your first line of defense



-  - Maintain SBOM for all services
-  - Track and update vulnerable dependencies regularly
-  - Auto-patch critical CVEs (don't wait for quarterly cycles)
-  - Use trusted sources for libraries/models only
-  - Lock dependency versions (no random latest pulls)
-  - Secrets management (no creds in repo, use vaults)
-  - Least privilege IAM everywhere
-  - Enforce SSO + MFA for all access
-  - Validate configs (no open buckets, public endpoints, etc.)

# Black-box testing (attacker mindset)

Now assume you're already exposed – try to break yourself



- Run regular pentests (internal + external)



- Fuzz APIs and inputs (especially AI prompts)



- Test for prompt injection / data exfiltration



- Abuse rate limits



- Check auth bypass scenarios



- Simulate real attacker flows



- Test exposed endpoints + misconfigs.

# Secure AI Interactions. Protect what matters

LLMs are powerful, but the real risk is data exposure.

Sensitive data should **never reach LLMs** or **go out to your users** via your systems.

**GUARDRAILS ARE MANDATORY WHEREVER LLM INTERACTIONS EXIST**



**Sensitive data must never reach LLMs or users via your systems.**

- Secrets & API Keys
- PII / Customer Data
- Internal Data
- Internal Code / Proprietary Info
- ... and more

**WHY IT MATTERS**

- Non-human identities and unmonitored access are high risk.
- Data can leak through prompts, responses and integrations.
- Exposed data can't be undone—externally or internally.
- Trust is lost when sensitive data reaches the wrong place.

**Discover and govern non-human identities. Apply guardrails everywhere. Keep your data and secrets protected—always.**

# SAST (code-level security)

Catch issues before runtime

`$ sast scan ./src`

- ✗ Hardcoded secret detected
- ✗ SQL Injection risk
- ⚠ Weak auth logic
- ✓ Input validation found

Scan complete.

```
1 function login(user, pass) {
2   const query = "SELECT * FROM users
3     WHERE user = '" + user + "'";
4     AND pass = '" + pass + "'";
5   return db.query(query);
6 }
7
8 const token = getToken();
9 if (token == null) {
10  allowAccess();
11 }
```

CI / CD

✓ ✓ ✓ ✗

SAST Scan

Issues Found: 3

- ✓ No hardcoded secrets
- ✓ Input validated
- ✓ SQLi protection
- ✓ Auth checks

Score: 9.2 / 10



- Run SAST on local code and in CI/CD



- Scan for hardcoded secrets



- Catch injection issues (SQL, command, etc.)



- Validate input sanitization paths



- Check auth logic (broken access control)



- Use linters + security rulesets



- Enforce secure coding standards



- Do manual code reviews for critical flows.

# Keep OS and apps patched and updated



**SYSTEM UPDATES**

- Windows OS  
Up to date
- Linux OS  
Up to date
- macOS  
Up to date

**APPLICATIONS**

- Google Chrome  
Up to date
- Slack  
Up to date
- VS Code  
Up to date
- Adobe Acrobat  
Up to date

- Closes security vulnerabilities
- Prevents known exploits
- Improves system stability and performance
- Ensures compatibility and reliability
- Reduces risk and downtime

Stay updated. Stay secure.

# Encrypt sensitive data



**Protect data. Preserve trust.**



**Protects data at rest**  
Ensures stolen data remains unreadable.



**Protects data in transit**  
Prevents interception and eavesdropping.



**Meets compliance**  
Helps satisfy security and regulatory requirements.



**Builds customer trust**  
Demonstrates commitment to data security.



**Reduces breach impact**  
Limits damage and speeds up recovery.

# Context aware access control



Right access. Right context. **Lower risk.**



Evaluate user +  
device + location



Grant access  
based on risk



Enforce adaptive  
authentication



Restrict access  
dynamically



Default to  
least privilege

# Periodic access reviews



## Review regularly

Validate who has access.



## Remove unnecessary access

Reduce risk and exposure.



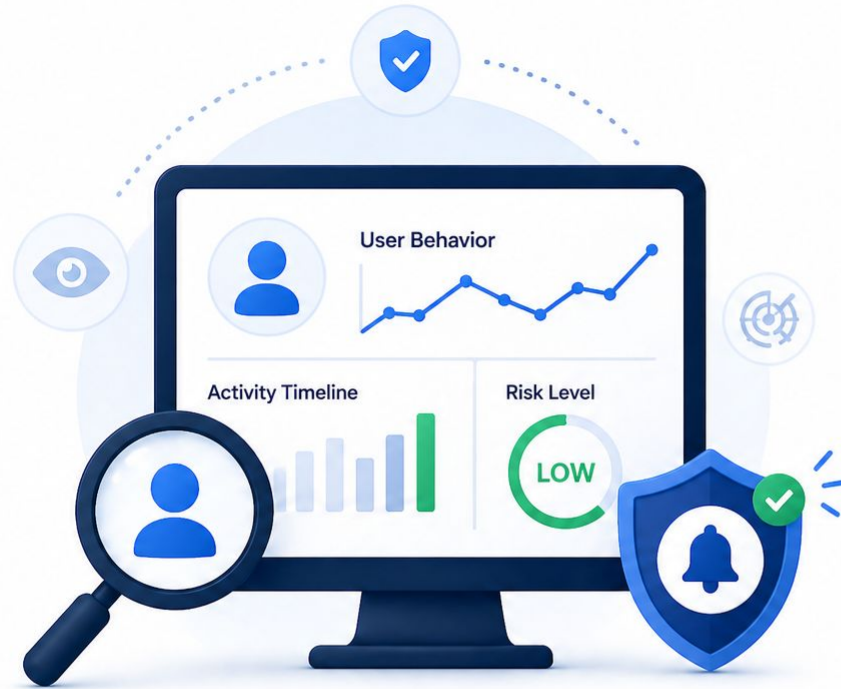
## Least privilege always

Access only what's needed.



Review access. Reduce risk.

# Behavioral monitoring



Baseline normal behavior



Detect anomalies in real time



Monitor privileged actions



Alert on suspicious activity



Automate response



Detect. Alert. Act. **Stay ahead.**

# Controlling AI Risk – Visibility, Governance, Guardrails

	AGENT TRAFFIC DETECTION	GOVERNANCE	LLM GUARDRAILS
<b>What it means</b>	Identify and monitor AI agents interacting with your systems.	Define and enforce rules for how AI is used in your org.	Runtime protections that filter, validate, and constrain AI inputs and outputs.
<b>Why it matters</b>	AI agents can scale attacks or misuse rapidly. Visibility first.	Prevent data leakage, ensure compliance, and avoid shadow AI.	Stops prompt injection, prevents data exfiltration, and keeps outputs aligned with policy.
<b>How it's done</b>	Behavior analysis, fingerprints, API gateways, and correlation.	Access control, policies, logging, audit trails, and approval workflows.	Input/output filtering, context isolation, policy rules, and retrieval constraints.
<b>Think</b>	<b>Think:</b> Who/what is actually calling my system?	<b>Think:</b> Are we using AI responsibly and within boundaries?	<b>Think:</b> Even if something goes wrong, the AI stays within safe boundaries.



**SEE**  
Visibility



**CONTROL**  
Policy & Governance



**PROTECT**  
Runtime Safety



**PRODUCT COMPANY OF THE YEAR  
PRIVACY PRODUCT OF THE YEAR**

**“जन गण मन बजवा देंगे”**



**Anirban Mukherji**

Founder & CEO, miniOrange

[anirban@miniOrange.com](mailto:anirban@miniOrange.com)