



SentryArk

SENTRYARK RESEARCH

Threat Report 2026

What we found across penetration tests, cloud reviews
and red team engagements.

■ CONTENTS

	Foreword	02
01	Key findings	03
02	What we found in testing	04
03	Red team lessons	06
04	Case studies	07
05	Attack techniques	09
06	Regulation in India	10
07	Recommendations	12
08	Research	14
09	About this data	14
	About SentryArk	15

■ FOREWORD

The most serious weaknesses we found this year were not exotic. They were applications that check who you are, but not what you are allowed to do.

This report brings together what our team found across penetration tests, cloud reviews and red team engagements up to September 2026. It covers 241 rated findings, the techniques behind them and the fixes that worked.

We wrote it for security leaders and engineering teams who want to know where real attackers would get in and what to fix first.

Every client is anonymised. No organisation is named, and no finding is tied to a specific company.

20

ENGAGEMENTS

241

RATED FINDINGS

81

CRITICAL AND HIGH

■ CHAPTER 01

Key findings

14 of 20

engagements found at least one critical or high issue.

Nine found a critical one.

5 of 9

web applications let a low-privilege user make themselves an administrator.

Usually one API call was enough, such as editing their own role field or calling a user-management function.

4 of 9

web applications let one customer reach another customer's data.

That included files, user accounts and, in one case, a customer database holding hundreds of thousands of records.

29 of 81

critical and high findings were on APIs, more than twice as many as the web interfaces in front of them (13).

A screen that hides data proves nothing about the API behind it.

83%

of critical and high findings (67 of 81) were found by manual testing.

Only 5 came plainly from a scanner.

18

critical and high findings came from two cloud reviews, more than a fifth of the total.

Over-privileged CI/CD pipeline identities were the most serious risk in both.

3

restricted areas, the data centre among them, opened to cloned access cards in our red team engagement.

Nobody noticed the network tap we planted there. The external perimeter held.

- Where clients retested, fixes held: all 19 retested critical and high findings were confirmed closed. Another 62 were never retested.


■ CHAPTER 02

What we found in testing

One class of flaw dominated: 22 of the 81 critical and high findings were authorisation failures, more than the next two categories combined.

Critical and high findings by type

81 in total



ASSESSMENT TYPE	ENGAGEMENTS	CRITICAL	HIGH	MEDIUM	LOW	INFO
Web application and API	9	8	22	25	27	7
External network	5	6	7	34	28	1
Internal network and endpoint	3	2	6	16	2	0
Cloud	2	7	11	13	1	0
Red team	1	6	6	3	3	0

Findings by assessment type. 20 engagements up to September 2026; categories grouped for readability.

Web applications and APIs

9 ENGAGEMENTS · 30 CRITICAL AND HIGH

The same mistake recurred across unrelated platforms: the server checked that a request came from a logged-in user, then trusted it. Users could promote themselves, create accounts in other tenants, reset other users' passwords and mark unpaid payments as paid.

Two newer weak spots stood out. One platform's password-reset code never expired and could be guessed. Another platform's AI helpers followed injected instructions in every feature.

Networks

8 ENGAGEMENTS · 21 CRITICAL AND HIGH

Most external perimeters showed only TLS and header issues. The serious exposure sat in portals built by third parties: an unauthenticated upload leading to code execution, SQL injection, and an HR API that accepted a fixed string in place of a login.

Internally, we found end-of-life database servers, obsolete SSL and every user running as a local administrator.

Cloud and identity

2 ENGAGEMENTS · 18 CRITICAL AND HIGH

Deployment pipelines were the identities to protect. In one Kubernetes estate every CI/CD identity held full cluster-admin, and one could grant any role across the cloud subscription.

Public storage, and permissions reaching from the cloud into Microsoft 365, made up most of the rest.

Hygiene

85 FINDINGS · 15 CRITICAL AND HIGH

Misconfiguration and weak cryptography made up 35% of all findings, but few were critical or high.

Most were old TLS versions, missing security headers and default settings that a configuration baseline would fix cheaply.



A login token proves who someone is. It does not prove what they are allowed to change.

■ CHAPTER 03

Red team lessons

The outside of the organisation held. The inside did not: cloned access cards and one default password took us from the front door to the data centre and to business credentials.

TARGET	SCOPE	OUTCOME
A large enterprise	External, physical and assumed breach	Most objectives reached; 12 critical and high findings

■ Perimeter HELD

VPN, Microsoft 365, Wi-Fi and VoIP resisted attack, and domain credential dumping was blocked.

■ Physical entry BYPASSED

Cloned RFID cards opened the main entrance, the data centre and the electrical room. Server racks were locked only at the front, so we attached a network tap from the back and joined it to the network as a new device. Nobody noticed.

■ Internal servers BYPASSED

An internal application server's admin console still used the vendor default password, which gave us code execution. From there we created a local administrator, installed remote-access software that worked over the internet, and found database and mail credentials on open file shares.

■ Endpoint security BYPASSED

Endpoint protection removed our first payload, but a modified one got through. A shadow-copy technique stole the password hashes it was meant to protect, and a malicious USB device ran commands on a managed laptop.

The engagement report recorded no time to objective and no SOC detection count, so this edition cannot measure detection. We will report detection rates from next year's engagements.

■ CHAPTER 04

Case studies

Five findings from five different organisations show what these numbers look like in practice. Each was confirmed by hand.

01 From the lowest account to platform owner

MULTI-TENANT B2B PLATFORM

On a multi-tenant B2B platform, the least-privileged user account sent one update to its own user record and raised its access level to super-admin. Two more flaws completed the chain. A password-reset function would reset any user's password, given only their ID, including users in other customer organisations, and a user-creation function made accounts at any level, in any tenant. The client fixed all three, and our retest confirmed they were closed.

- **Lesson.** A login token proves who someone is, not what they may change. Fields such as role and status must never be writable by the user.

02 One pipeline identity, the whole cloud

SAAS COMPANY, KUBERNETES ESTATE

In a read-only review of a SaaS company's Kubernetes estate, we listed the cluster's role bindings by hand. Every application's CI/CD identity held full cluster-admin, and several were trusted by more than a dozen code repositories each. One could also grant any role across the cloud subscription. The automated scanner reported nothing, because it looks for risky custom roles, not for many identities bound to the built-in admin role.

- **Lesson.** In modern cloud estates, the deployment pipeline is the identity attackers want. Scope each pipeline to its own namespace and block new cluster-admin bindings.

83%

of critical and high findings were found by manual testing.

67 of 81. Only 5 came plainly from a scanner.

03 AI features that take orders from users

BUSINESS APPLICATION WITH AI HELPERS

A business application had added AI helpers that draft text for official records. Typing variations of "ignore previous instructions" made every AI feature drop its task, reveal details about the model and produce harmful content. That output landed in the same fields an employee would fill in, where managers would later read it.

- **Lesson.** Every AI feature failed the same way, which points to how the model calls were built rather than to one bad prompt. Separate instructions from user input, and moderate the output.

04 A supplier portal open to code execution

LARGE ENTERPRISE, EXTERNAL TEST

During an external test of a large enterprise, we found an upload API on a supplier portal that needed no login and built the save location from fields in the request. Path traversal let us place files outside the upload folder, including one that would run code when the service restarted. The same backend was also open to SQL injection against its production database.

- **Lesson.** The perimeter was mostly clean. The serious exposure sat in portals built by third parties, which need the same testing as your own code.

05 MFA does not stop a stolen token

ENDPOINT AND MICROSOFT 365 REVIEW

In an endpoint and Microsoft 365 review, we took a standard user's laptop, extracted its cloud sign-in token with a public tool and replayed it from another machine. The token opened the organisation's directory, with no MFA prompt at any point. Every user was a local administrator, and no EDR was installed, so nothing stopped or flagged the theft.

- **Lesson.** MFA protects the sign-in, not the session that follows. EDR with token-theft detection, no local admin rights and conditional access close the gap.

■ CHAPTER 05

Attack techniques behind our findings

The technique behind the most critical and high findings was account manipulation: changing roles, permissions and passwords that an attacker should never reach.

We mapped 143 of the 241 findings to a MITRE ATT&CK technique. They describe how each issue could be exploited, not activity seen in a live attack. This edition has no SOC incident data, so the last column is what we recommend a SOC watch for.

TECHNIQUE	WHAT IT LOOKED LIKE IN OUR TESTS	CRITICAL AND HIGH	ENGAGEMENTS	WHAT A SOC SHOULD WATCH FOR
T1098 Account Manipulation	Users promoting themselves, rewriting role permissions, resetting other users' passwords	13	7	Role and permission changes made by non-admins or by users on themselves; log authorisation decisions in the application
T1190 Exploit Public-Facing Application	SQL injection, path traversal and unauthenticated uploads on internet-facing portals	10	3	Injection and traversal patterns in web logs, new files written to web directories, an inventory of third-party portals
T1078 Valid Accounts	Default passwords, over-privileged pipeline identities, every user a local admin	6	6	Logins to default accounts, service identities used outside their pipelines, admin rights granted outside provisioning
T1530 Data from Cloud Storage	Public buckets and blob containers, cross-tenant file listing and download	6	5	Anonymous storage access, bulk downloads, listing outside a tenant's own path
T1552 Unsecured Credentials	Cluster secrets readable by app roles, passwords on open file shares	4	4	Secret and vault reads by unusual identities, credential files on shares
T1136 Create Account	Accounts created at any level or in any tenant, new local administrators	3	4	New directory or local admin accounts created outside the normal process
T1528 Steal Application Access Token	Sign-in tokens taken from a laptop, APIs returning other users' tokens	2	3	Token replay from a new device or network, refresh tokens reused after logout

■ CHAPTER 06

Regulation in India

The date that matters most is May 2027, when the DPDP Rules' security and breach-notification duties take effect. Banks and payment firms already face stricter RBI rules that have applied since 31 July 2026.

RULE	APPLIES TO	WHAT IT ASKS FOR ON SECURITY	WHEN
DPDP Rules, 2025	Any organisation processing digital personal data in India	Encryption or masking, access control and access logs, detection of unauthorised access, backups, and logs kept for at least one year. Breaches go to the Data Protection Board without delay, with a detailed report within 72 hours, and to affected people.	Notified 13 Nov 2025. Consent managers from Nov 2026; most duties from May 2027
CERT-In Directions	Service providers, intermediaries, data centres, companies and government bodies	Report cyber incidents within 6 hours of becoming aware. Keep logs for 180 days in India. Sync clocks to NIC or NPL time servers.	In force since 2022
RBI technology Master Directions	RBI-regulated entities	Report incidents within 6 hours through RBI's DAKSH platform. Vulnerability assessments every 6 months and penetration tests every 12 months for critical and internet-facing systems. Board oversight of technology risk.	Issued and in force 31 Jul 2026
SEBI CSCRF	SEBI-regulated entities	Security monitoring through a SOC (own, group, market or third-party), cyber audits, and VAPT after every major release.	Issued Aug 2024; most entities had to comply by 31 Aug 2025
IRDAI Cybersecurity Guidelines, 2023	Insurers and insurance intermediaries	Security controls, incident response plans, regular security audits, a named C-suite owner and control over third parties.	Compliance due by Apr 2024

- Under the DPDP Act, failing to take reasonable security safeguards can draw a penalty of up to ₹250 crore.

Key dates

APR 2022	☐	CERT-In Directions 6-hour incident reporting and 180-day logs
APR 2024	☐	IRDAI guidelines Compliance due for insurers and intermediaries
AUG 2025	☐	SEBI CSCRf Compliance due for most SEBI-regulated entities
NOV 2025	☐	DPDP Rules notified Phased start of India's data protection rules
JUL 2026	☐	RBI Master Directions In force on issue for RBI-regulated entities
NOV 2026	☐	DPDP consent managers Consent manager rules take effect
MAY 2027	☒	DPDP security and breach duties Safeguards, logs and 72-hour breach reports apply

What to do before 2027

Cross-tenant access is a breach waiting to happen.

Four of nine web applications let one customer reach another's data. Under DPDP, that is a notifiable breach with a 72-hour report.

You cannot report in 6 hours what you cannot see.

Our red team went undetected through a data centre and onto servers. Make sure logs and alerting reach the systems that hold personal data.

Keep the right logs for the right time.

CERT-In expects 180 days of logs held in India. DPDP adds at least one year for logs tied to personal data.

Put third parties in scope.

Our most serious external findings were in vendor-built portals, and every one of these rules makes you answerable for your suppliers.

As of October 2026. This chapter summarises the rules and is not legal advice.

Sources: AZB & Partners and KPMG India on the DPDP Rules; Lawrbit on the CERT-In Directions; KPMG India on the RBI Master Directions; Taxmann and Uniqus on SEBI CSCRf; King Stubb & Kasiva on the IRDAI guidelines.

■ CHAPTER 07

Recommendations

Ten actions would have prevented most of the critical and high findings in this report. Each comes from a pattern we saw in at least one engagement, with a first step for each size of organisation.

01 Enforce authorisation in one shared place

Based on: 5 of 9 web apps let users make themselves admin

STARTUPS

Use your framework's policy layer; never hand-roll checks per endpoint

MID-SIZE

Add regression tests that expect cross-tenant calls to fail

LARGE ENTERPRISES

Make it a platform service every team must use

02 Test APIs directly, with every role and tenant

Based on: 29 of 81 critical and high findings were on APIs

STARTUPS

Test before each major release

MID-SIZE

Test twice a year and after big changes

LARGE ENTERPRISES

Test continuously, plus an annual deep test

03 Never let users write role, status, tenant or payment fields

Based on: Users promoting themselves, and payments marked as paid from the browser

STARTUPS

Allow-list editable fields per role

MID-SIZE

Review every update endpoint

LARGE ENTERPRISES

Block these fields with a shared API gateway rule

04 Treat pipelines and service identities as your most privileged accounts

Based on: 18 critical and high findings in 2 cloud reviews

STARTUPS

One deploy identity per app, no cluster-admin

MID-SIZE

Scope by namespace and block new admin bindings

LARGE ENTERPRISES

Review service identities quarterly, like admins

05 Block public cloud storage at account level

Based on: Public buckets with CVs, contracts and pricing

STARTUPS

Turn on the account-wide block now

MID-SIZE

Alert on any new public access

LARGE ENTERPRISES

Enforce by policy across all subscriptions

06 Test third-party portals and vendor platforms

Based on: The worst external findings sat in vendor-built portals

STARTUPS

Ask vendors for recent test reports

MID-SIZE

Include vendor portals in your own test scope

LARGE ENTERPRISES

Write testing and fix times into contracts

07 Secure AI features before launch

Based on: Every AI feature in one app followed injected instructions

STARTUPS

Separate system instructions from user input

MID-SIZE

Add output moderation and rate limits

LARGE ENTERPRISES

Red team AI features like any other release

08 Back MFA with endpoint and session controls

Based on: A stolen sign-in token worked without any MFA prompt

STARTUPS

Remove local admin rights

MID-SIZE

Deploy EDR with token-theft detection

LARGE ENTERPRISES

Add conditional access and device-bound sessions

09 Test your doors and your endpoint tools

Based on: Cloned cards opened the data centre; EDR was bypassed

STARTUPS

Badge-only access to server rooms

MID-SIZE

Rear rack locks and alerts on new network devices

LARGE ENTERPRISES

Run physical red teams and tune EDR after each

10 Fix hygiene with a baseline, and retest every critical and high finding

Based on: 85 hygiene findings; 62 of 81 critical and high never retested

STARTUPS

Change default passwords and enforce TLS 1.2 or later

MID-SIZE

Apply a configuration baseline to every host template

LARGE ENTERPRISES

Track fix-and-retest times as a board metric

19/19

retested critical and high findings were confirmed closed.

Another 62 were never retested. A fix you have not checked is a fix you are hoping for.

■ CHAPTER 08

Research

Members of our technical team have published two CVEs, and our Team Yazhi researchers hunt for new threats and vulnerabilities between client engagements.

■ CHAPTER 09

About this data

This report draws on 20 SentryArk engagements delivered up to September 2026, which produced 241 rated findings: 29 critical, 52 high, 91 medium, 61 low and 8 informational.

WHAT WAS TESTED	9 web application and API assessments, 5 external network tests, 3 internal network and endpoint reviews, 2 cloud reviews (Azure, Kubernetes and Microsoft 365), and 1 red team engagement covering external, physical and assumed-breach testing.
WHO WAS TESTED	Mostly SaaS companies, plus three larger enterprises. Clients are never named, and no finding is tied to a specific organisation.
HOW	Severities are taken from our delivered reports. Categories and OWASP or MITRE ATT&CK references were assigned consistently so reports can be compared.
LIMITS	The sample is small, and two organisations account for 61% of findings, so read percentages as patterns rather than industry rates. Only 44 findings were retested. Red team results come from a single engagement, and this edition contains no SOC or dedicated Active Directory data.

ABOUT SENTRYARK

We test like attackers, so you can fix what they would find.

SentryArk is an information security and cyber resiliency provider. Our consultants hold OSCE, CRT0, CEH and ISMS certifications, and members of our team have published 2 CVEs. We have secured 20+ organisations, assessed 8K+ devices and 30+ applications, and reported 250+ vulnerabilities.

VAPT · APPLICATION AND API SECURITY · MOBILE SECURITY · CLOUD SECURITY · ACTIVE DIRECTORY · RED TEAMING · 24/7 MANAGED SOC · SIEM, SOAR AND EDR · RISK ASSESSMENT · ISO 27001 · DEVSECOPS

TALK TO US

To discuss any finding in this report, or to test your own environment:

contact@sentryark.com

www.sentryark.com



SentryArk