CMWAP · WORKED EXAMPLE

Pentest Plan — RatCTF

ratctf.com in an 8-hour window. Five testable hours, and a written reason for every endpoint that does not get them.

> **Read this beside the template.** Same sections, same clock, filled in against a target far larger than the window. That is the point: the platform has roughly 96 challenge machines, eight multi-machine ranges and a portal, and this plan tests **eleven endpoints**. The other rows are not missing — they are marked **OUT** or **DEFERRED**, each with a reason and, where it applies, the condition that would pull it forward.

Field	Value
Candidate	A. Candidate
Target / engagement	ratctf.com — RatCTF training platform
Window opens (T+0:00)	Mon 09:00 CET
Plan due (T+0:45)	**Mon 09:45 CET** — submitted 09:41
Report + debrief due (T+8:00)	Mon 17:00 CET
Authorization	Engagement letter dated 2026-08-07. The platform rules at `/rules` apply as the baseline and are not overridden by this plan. Disclosure route: `/.well-known/security.txt`.
Plan version	v1.0

1 · Objective

0:00 – 0:10

Determine whether a free-tier RatCTF account can obtain content or standing it has not earned — premium or retired machines, another user's progress, or leaderboard position from challenges it never solved — and state what that would cost RatCTF in subscription revenue and leaderboard credibility.

One question, answerable in five hours. Not “assess the security of RatCTF”, which is a six-week engagement and would have produced a plan I could not deliver.

2 · Time budget

0:10 – 0:15

Fixed cost — comes off the top	Hours
Total window	8.00
This plan	−0.75
Report — ring-fenced, never borrowed from	−1.50
Debrief video: record, review, upload	−0.50
Setup, access, evidence admin	−0.25
ACTUAL TESTABLE TIME	**5.00 h = 300 min**

Wall clock	Block	Budget	What must be true when it ends

09:00–09:45	Plan — write and send	45 min	Submitted at 09:41
09:45–10:45	Recon and mapping	60 min	Two free accounts and the premium account live; Burp proxying; endpoint table confirmed against the real app
10:45–13:00	IN · MUST — rows 1–5	135 min	Entitlement and scoring questions answered either way
13:00–14:15	IN · SHOULD — rows 6–11	75 min	Supporting rows tested or explicitly parked
14:15–14:45	Verify and capture proof	30 min	Every finding re-run from its own written steps on a clean session
14:45–16:15	Report	90 min	Ring-fenced
16:15–16:45	Debrief video	30 min	Recorded and uploaded
16:45–17:00	Submit and buffer	15 min	Both deliverables in
Testing subtotal		**300 min**	**= 5.00 h**

3 · Endpoint inventory and disposition

0:15 – 0:35

IN — authorized, relevant, affordable. **OUT** — not authorized. Never tested. **DEFERRED** — authorized and relevant, does not fit 300 minutes.

#	Endpoint / function	Method	Role(s)	WSTG ID(s)	Disp.	Why this disposition	Budget
1	`/challenges/{slug}` — premium / retired gate	GET	free	WSTG-ATHZ-02 WSTG-ATHZ-04	**IN**	This gate *is* the subscription revenue named in the objective. Direct object reference on a slug the account has no entitlement for.	40 min
2	`/api/flags/submit` — flag submission	POST	free ×2	WSTG-BUSL-01 WSTG-BUSL-02	**IN**	Converts work into leaderboard standing. Tests whether a flag is accepted for a machine the account never spawned.	35 min
3	`/api/instances/{id}` — spawn / reset / destroy	POST/ DELETE	free ×2	WSTG-ATHZ-04 WSTG-BUSL-02	**IN**	Instance ownership boundary. If account A can reset account B's instance, one learner can wreck another's engagement.	25 min
4	`/vaults/{id}/claim` — Treasure Vault claim	POST	free	WSTG-BUSL-02 WSTG-ATHZ-02	**IN**	A one-time, scarce, paid-tier action — the classic replay and double-claim target.	20 min
5	`/leaderboard` + scoring recalculation	GET	free	WSTG-BUSL-01 WSTG-ATHZ-04	**IN**	Second half of the objective. Checks whether standing can be inflated without solving.	15 min
6	`/profile` — display name, avatar, bio	POST	free	WSTG-INPV-02 WSTG-ATHZ-03	**IN**	Stored XSS renders on the leaderboard, seen by every user including staff. Also the likeliest mass-assignment surface (role, tier).	20 min
7	`/paths`, `/quests` — progression unlocks	GET/ POST	free	WSTG-BUSL-02	**IN**	Workflow with an ordering assumption; skipping a step is cheap to test and feeds row 5.	15 min

8	`/api/docs` — published API surface	GET	free	WSTG-CONF-05 WSTG-INFO-02	**IN**	Read-only enumeration to confirm rows 2–4 are complete. No requests fired from it beyond those rows.	10 min
9	`?next=` / `?redirect=` params on nav routes	GET	free	WSTG-CLNT-04	**IN**	Cheap, and an open redirect chains into row 1 by carrying an entitled session somewhere it should not go.	10 min
10	`/premium` — entitlement display vs enforcement	GET	free, premium	WSTG-ATHZ-02	**IN**	Paired with row 1: is the gate enforced server-side or only rendered? 10 minutes with two sessions answers it.	10 min
11	Set-Cookie on session issue and on tier upgrade	—	free, premium	WSTG-SESS-02 WSTG-SESS-03	**IN**	Two minutes of observation; if the session does not rotate on upgrade it undermines rows 1 and 10.	10 min
IN subtotal — matches the 210-minute MUST + SHOULD budget							**210 min**
12	96 single challenge machines (Injectrix, Loophole, Debugtrap, …)	—	free	full INPV / ATHZ / BUSL set	**DEFERRED**	Authorized and in-theme, but 96 machines at even 20 min each is 32 hours. Cannot be sampled honestly in 300 minutes.	—
13	CMWAP Range — turnstile, nightglass, crossfire, ledgerloop, cacheleak	—	free	WSTG-ATHZ-02, INPV-02, SESS-05, BUSL-01, CONF-04	**DEFERRED**	The single most relevant range to the objective's classes, but 5 machines needs ~2 h and would take the whole MUST block.	—
14	Pentest Range — IronClad, MedBridge, FinTrust	—	free	WSTG-ATHZ-01, BUSL-01	**DEFERRED**	Each is itself a 24-hour engagement. Including one would consume the window and answer a different question.	—
15	AD / network ranges — CorpNet, VaultNet, MegaCorp, OperationSwitchboard	—	free	outside WSTG scope	**DEFERRED**	Network and AD classes, not web. Authorized, but outside the objective and outside the CMWAP class set.	—
16	`/hub`, `/tips`, `/reviews` — community content	GET/ POST	free	WSTG-INPV-02, CLNT-01	**DEFERRED**	Real stored-XSS surface, but row 6 already covers the same sink class on a higher-value page.	—
17	“The Burrow” — undocumented area	GET	free	WSTG-CONF-05	**DEFERRED**	Interesting and exactly the kind of thing that eats an afternoon. Named here so that if I test it, it is a logged decision and not drift.	—
18	`/login`, `/signup`, session handling	POST	—	WSTG-ATHN-03, IDNT-04	**OUT**	Portal authentication. Excluded by `/rules`; not in the engagement letter. Would need separate written authorization.	—

19	`/api/*` portal endpoints as infrastructure	—	—	—	**OUT**	Excluded by `/rules`. Rows 2–4 exercise them only as a normal user would, never as scan targets.	—
20	Kubernetes, container orchestration, cloud accounts	—	—	WSTG-CONF-11	**OUT**	Explicitly excluded by `/rules` and outside the engagement letter's asset list.	—
21	Other users' accounts, sessions, vault claims	—	—	WSTG-ATHZ-04	**OUT**	Third-party data. No lawful basis. Row 3 uses two accounts I own for the same question.	—
22	Checkout and payment processor	POST	—	WSTG-BUSL-01	**OUT**	Third-party system. Testing it would test someone else's asset.	—
23	Automated scanning / DoS against any portal endpoint	—	—	—	**OUT**	Explicitly forbidden by `/rules`, and destructive to other learners mid-engagement.	—

The line I will not cross. If an IN row reveals a route out of a challenge container, into orchestration, or into real user data, I stop, do not develop it further, and escalate to `security@ratctf.com` within the hour. Confirming reachability is the finding; proving impact would put me outside my own table.

4 · Deferred register

included above

Deferred item	WSTG ID(s)	Why it did not fit	What would pull it forward
CMWAP Range (5 machines)	ATHZ-02, INPV-02, SESS-05, BUSL-01	~2 h; would consume the entire MUST block	If rows 1–5 close before 12:15, take turnstile only (~25 min) as it overlaps row 1's question
Community content — hub, tips, reviews	INPV-02, CLNT-01	Duplicate sink class to row 6, lower value	If row 6 finds stored XSS, spend 15 min confirming the same sink here to widen impact
“The Burrow”	CONF-05	Unknown size — cannot be budgeted honestly	Only in the verify block, only if everything else is closed, capped at 15 min and logged
96 single machines + Pentest Range + AD ranges	full set	32 h+ of work against a 5 h budget	Nothing in this window. Recommended as a separate engagement in the report.

5 · Rules of engagement

0:35 – 0:40

Environment	Production — the only one that exists. No destructive tests, no persistence.
Accounts	Two free accounts I create (named in the report) + one premium account supplied by the client
Rate limit	Max 10 concurrent requests to any challenge host, 100 ms delay. No automated tooling pointed at portal endpoints at any rate.
Will NOT do	No DoS, no destructive writes, no persistence on shared machines, no interaction with other users' accounts or vault claims, no publishing of flags
Out-of-band listener	interactsh, stood up in the recon block; host recorded in the notes vault
Evidence	Request, payload, response and a one-frame proof shot per finding, stored locally; nothing uploaded to third parties

Escalation	Boundary crossing, real user data, or availability impact → stop, notify `security@ratctf.com` within the hour, log it
Abort criteria	Signs of a third-party intrusion; evidence my testing degraded the platform for others; any finding I cannot explore without leaving this table
Tools declared	Burp Community, ffuf (10 threads), jwt_tool, interactsh, Obsidian — all free. AI assistant (free tier) for payload variants and a first report draft; 20 min budgeted inside the report block to verify everything it produces by hand.

6 · Thesis, risks and change log

0:40 – 0:45

Thesis

The serious finding is in **row 1**. A platform that gates premium and retired content by slug almost always checks entitlement in the page that lists machines rather than in the handler that serves one. If that holds, a free account reads paid content by direct reference — which is the revenue question the objective asks, and it chains into row 10. I test row 1 first, at 10:45.

Assumptions and risks

Assumption or risk	What I do if it turns out false
The supplied premium account reflects a real paid entitlement	Row 10 becomes untestable as designed; note the substitution in the change log and say so in the debrief
Challenge instances are per-user, not shared	Row 3 is dropped — it cannot be tested without touching another user. Recorded as untested, not as clean.
Recon confirms the endpoint names in this table	Rows are renamed in the change log; dispositions and budgets carry over unchanged
210 minutes of IN work is achievable at the stated rate limit	Rows 9 and 11 are the first to go — 20 min recovered, both are supporting evidence rather than the objective

Change log

Empty at submission. These three were added during the window — this is the section the examiner reads to check adherence.

T+	What changed	Cost	Paid for by
T+2:10	Row 1 overran its 40-minute box — the entitlement bypass looked reachable and was worth finishing	+25 min	Row 9 (open redirect) dropped; recorded as untested
T+4:05	Row 3 abandoned — could not spawn a second concurrent instance reliably, and testing it any other way needs another user's account	−25 min	Returned to the SHOULD block; row 7 got the time
T+4:50	Pulled turnstile forward from the deferred register per its stated trigger, capped at 25 min	+25 min	The recovered row 3 time. Trigger condition was met and is quoted in the report.

Why this plan scores

Planning — 210 IN minutes against a 210-minute budget; the arithmetic closes. **Adherence** — every finding in the report traces to an IN row, and all three deviations are logged with what paid for them. **Findings** — impact stated in revenue and leaderboard terms, matching the objective. **Time spent** — planned versus actual is readable straight off the wall-clock table. **Next iteration** — row 3 was the wrong row to plan; it needed an assumption I could not verify before committing to it.

The XSS Rat · CMWAP · Worked example for teaching. Authorized testing only — RatCTF's published rules at `/rules` govern that platform and this plan does not override them.