THE CYBER TOOLCHAIN

NO. 102

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 11, 2026 · EVERY DAY

The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 24 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 7 tools

▾ // THE BRIEF HIDE

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out. adscan now actively confirms ADCS ESC9/ESC10/ESC15 rather than flagging them as theoretical. Harden-er extends egress deny-listing to macOS runners and lands on

vim ? help

GitHub Enterprise Server. And Google SecOps lets analysts write SIEM searches in GoogleSQL or plain language, with data RBAC enforced on results.

PROTECT

Prove an ADCS escalation path is real before you write it up

adscan adds active confirmation for ESC9/ESC10/ESC15, so a misconfigured certificate template stops being a maybe-finding that a domain admin can argue with. sara's audit engine now runs over the RouterOS API with 22 structured checks plus a CVE lookup, which removes the SSH scraping that made MikroTik auditing fragile.

adscan · sara

PROTECT

Stop a compromised build step from phoning out, on macOS and on self-hosted Enterprise

Harden-Runner's deny-list egress reaches macOS runners and GitHub Enterprise Server, closing the two gaps where most teams had no network control over CI at all; namespace-scoped RBAC in ARC limits what one team's runners can see on a shared cluster. Socket's reachability engines add JVM classpath integration through the CLI sidecar, so a Java advisory can be judged on whether the vulnerable code path is actually loaded.

StepSecurity · Socket

DETECT

Query the SIEM without knowing the query language

SecOps Search takes GoogleSQL and Gemini-generated natural language, with UDM Lookup and data RBAC applied to what comes back — a junior analyst can ask a question without an hour of syntax, and the access boundary holds. Censys AI Investigations pivots across indicators automatically and ranks what it finds, collapsing the manual hop-by-hop infrastructure walk.

Google SecOps · Censys

OFFENSIVE SECURITY

◆ Exploitation & C2

adscan

Feedback ⓘ v13.2.0 NOTES ↗ RANK

Free Active Directory pentesting tool for Linux, macOS and Windows. Automates AD and LDAP enumeration, Kerberoasting, ASREPRoast, password spraying, ADCS/ESC1, DCSync, RBCD, gMSA, shadow credentials, NTLM relay and credential dumping across 104 techniques, mapping BloodHound-compatible attack paths to Domain Admin.

> adscan v13.2.0 adds ENS compliance-scope scoring, GPU-accelerated local hash cracking in CI, fully branded/white-label report deliverables, headless credential-dumping subcommands, and active confirmation of ADCS ESC9/ESC10/ESC15, alongside several breaking changes to report defaults.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 **Branded and white-label report deliverables** NEW 90

Adds `adscan deliver --branding {adscan,cobrand,whitelabel}` together with `--client-cover`, `--client-back`, `--client-logo`, `--theme <file.css>`, and `--emit data` (writes `report_model.json`) to produce co-branded, white-label, or stylesheet-customised deliverables including an HTML edition of the Security Assessment Report.

Produce a white-label client report with your own cover, logo, and CSS stylesheet, plus a machine-readable data export for further processing.

```
$ adscan deliver --branding whitelabel --client-cover cover.pdf --client-back back.pdf --client-logo logo.png --theme client.css --emit data
```

— Full flag set with a runnable deliver command v13.2.0

02 **Explicit framework selection required for audit reports** 80

BREAKING

`adscan ci --generate-report` no longer defaults to ENS on audit runs; adds `--frameworks` as a required flag, and the tool refuses to render without an explicit framework choice.

Generate a compliance-scoped ENS audit report for a specific framework without the tool silently defaulting to the wrong regime.

```
$ adscan ci auth --type audit --interface eth0 --domain
corp.local --dc-ip 10.0.0.1 -u alice -p 'S3cr3t!' --generate-
report --frameworks ENS
```

— Breaking flag change shown with a runnable command v13.2.0

03 **ENS compliance scope declaration** NEW 75

Adds `adscan execute compliance_scope` to declare an ENS scope (profile plus per-dimension security levels) so findings are scored against the client's actual regime rather than an assumed category; CCN-STIC 890A and 890C profiles use their own applicability tables.

— Command and profiles named, no runnable example given v13.2.0

04 **Active confirmation of ADCS ESC9, ESC10 and ESC15** IMPROVED 75

Adds `ADSCAN_NO_ESC15_CONFIRM=1` and `ADSCAN_NO_ESC9_CONFIRM=1` to skip the self-subject certificate issuance that now actively confirms ADCS ESC15, ESC9, and ESC10 findings instead of inferring them from template attributes.

— Mechanism and opt-out env vars named, no example v13.2.0

05 **GPU-accelerated local hash cracking in CI** IMPROVED 70

Adds `ADSCAN_NO_LOCAL_CRACKING=1` to opt out of local NT hash cracking; GPU cracking now works in `adscan ci` and `adscan execute`, not only in the interactive workbench, and runs against a time budget measured on your hardware.

— Env var and commands named, no worked example v13.2.0

06 **Full credential table and explicit writeup generation** 70

BREAKING

Adds `--full-credentials` to `adscan writeup` to include the full recovered-credential table, replacing the default audit output which shows secrets only by their shape; the lab writeup is also no longer produced automatically on any audit path including `adscan ci`, so it must now be run explicitly on an audit workspace.

— Flag and migration note named, no runnable example v13.2.0

07 **Expanded Client Deliverable Kit reports** NEW 70

Adds `endpoint_defense_appendix.csv` to the Client Deliverable Kit, listing per-host antivirus and EDR posture, the credential-extraction techniques attempted, and their outcomes, and adds an AD Password Hygiene Report whenever credential recovery

ran, covering recovery rate, coverage per privilege tier, and groups of accounts sharing a password.

— Named deliverable files and report contents, no example `v13.2.0`

---

08 **Headless credential dumping subcommands** NEW 65

Adds `dump_lsa`, `dump_sam`, `dump_lsass`, `dump_dpapi`, `dump_registries`, and `dump_host` as subcommands of `adscan execute`, each accepting its own target and credentials, so credential dumps can run without the interactive workbench.

— Six subcommands named, no invocation example shown `v13.2.0`

THINNER COVERAGE BELOW

09 **CVSS score display control in client reports** BREAKING 55

CVSS scores are now removed from client deliverables by default; adds `ADSCAN_REPORT_CVSS_BASE=1` to print the CVSS Base score beside its vector for engagements that require it by contract.

— Breaking default and env var named, no example `v13.2.0`

---

10 **LDAP channel binding assessment over StartTLS** IMPROVED 55

Adds LDAP channel binding assessment over StartTLS when a domain controller firewalls LDAPS, so enforcement is now reported as Required or Not enforced instead of Unknown.

— Describes behaviour change, no command or example `v13.2.0`

└──▷ BREAKING ON UPGRADE

! `adscan ci --generate-report` no longer defaults to ENS on audit runs — it now requires `--frameworks` and refuses to render without it.

! Client deliverables no longer include CVSS scores by default; set `ADSCAN_REPORT_CVSS_BASE=1` to restore the Base score beside its vector.

! The lab writeup is no longer produced on any automatic audit path including `adscan ci`; run `adscan writeup` explicitly on an audit workspace to generate it.

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

sara

Feedback ⓘ v1.4.0 NOTES ↗ RANK

RouterOS Security Assessment Toolkit

> Sara v1.4.0 rewrites its audit engine to run over the RouterOS API with 22 structured checks and adds a new CVE lookup command, while dropping SSH-based connectivity entirely.

└──▷ **WHAT SHIPPED · 3 FEATURES** most completely described first ? what's the number?

01 **RouterOS API audit engine with 22 checks** NEW 95

`sara audit <IP> <USERNAME>` now connects over the RouterOS API (`api-ssl` , port 8729 by default) to run 22 named checks (MGMT-001 through WIFI-002), each returning severity, live evidence pulled from the router, and a RouterOS fix command. The `--port` flag targets a non-standard API port, and `--plaintext` switches from `api-ssl` to the unencrypted `api` service (port 8728).

Audit a router that exposes the API on a non-standard port, e.g. in a hardened deployment where 8729 is remapped.

```
$ sara audit 192.168.88.1 sara --port 8730
```

— Full command, flags, ports, and check range all named. v1.4.0

02 **Transport migrated from SSH to RouterOS API** BREAKING 78

Sara's connection transport moved from SSH (via Netmiko) to the RouterOS API (via librouteros), dropping dependencies on `python3-netmiko` and `python3-packaging` . Routers must now have the `api-ssl` service (port 8729) enabled and an account with `read` and `api` policies; SSH-based connectivity no longer works.

— Names exact dependencies, port, and required account policies. v1.4.0

03 **CVE lookup command via NVD** NEW 75

`sara cve <IP> <USERNAME>` looks up CVEs for the RouterOS version running on a connected router via the NVD API. `sara cve --version <VERSION>` performs the same lookup for any RouterOS version without connecting to a router, useful for offline or pre-upgrade risk assessment.

Look up CVEs for a specific RouterOS version without connecting to a device — useful for offline or pre-upgrade risk assessment.

```
$ sara cve --version 7.14.3
```

— Exact commands and flag given; underlying NVD mechanism only briefly described. v1.4.0

└──▷ **BREAKING ON UPGRADE**

! Transport migrated from SSH (Netmiko) to the RouterOS API (librouteros); routers must now have the `api-ssl` service (port 8729) enabled and an account with `read` and `api` policies — SSH-based connectivity no longer works.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ

New controls and chat improvements in Copilot for JetBrains (2026-10-10) BLOG ↗

GitHub is a web-based platform for version control and collaborative software development using Git repositories.

Copilot for JetBrains adds enterprise default-model controls, an MCP startup toggle, an inline Fix action from diagnostics, and richer chat and account UX.

Adds a new MCP setting to disable automatic server startup for Copilot and Claude, giving administrators control over when configured tools become active.

Enterprise administrators can now select any available Copilot agent model as the organization-wide default for new conversations via managed settings.

Diagnostic intention menus now include a `Fix` action that opens inline chat in agent mode (or ask mode when agent is unavailable) to repair the reported issue.

Adds a `Switch Account` action to signed-in Copilot account controls, along with clearer labels, a simpler status display, and better navigation for long account lists.

Chat messages are now grouped by turn with controls for moving quickly between turns.

+3 more

RETIRED SUPPORTED PLATFORM JetBrains IDE 2025.1 support ended

└──▷ BREAKING ON UPGRADE

! GitHub Copilot has ended support for JetBrains IDE 2025.1; JetBrains IDE 2025.2 or later is now required to use the plugin.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

StepSecurity

Feedback ⓘ

Harden Runner Releases: October 9, 2026 (2026-10-09) BLOG ↗ RANK

StepSecurity hardens CI/CD pipelines by detecting and preventing supply chain attacks through runtime security monitoring and vulnerability detection.

> StepSecurity's Harden-Runner now supports GitHub Enterprise Server environments, adds a deploy option for ephemeral self-hosted VMs, extends deny-list egress to macOS, and brings namespace-scoped RBAC to ARC on Kubernetes.

└──▷ **WHAT SHIPPED · 4 FEATURES** most completely described first ? what's the number?

01 **Namespace-scoped RBAC for ARC Harden Runner** NEW 86

Adds `rbac.scope: namespaced` to ARC Harden Runner (v2.21.0) to install with only `Role` and `RoleBinding` objects in the release and runner namespaces, with no `ClusterRole` or `ClusterRoleBinding`, for multi-tenant Kubernetes clusters.

Install ARC Harden Runner on a multi-tenant Kubernetes cluster without cluster-scoped RBAC resources.

yaml

```yaml
rbac:
  scope: namespaced
```

— Config key, exact RBAC objects, and use case all named.

Harden Runner Releases: October 9, 2026…

02 **GitHub Enterprise Server support for Harden-Runner** NEW 83

Adds `customer` and `server-name` inputs to the `step-security/harden-runner` Action (v2.22.0) to enable Harden-Runner on GitHub Enterprise Server environments, required for TLS inspection, policy fetching, and agent deployment. The Harden-Runner agent (Self-Hosted VM v1.9.3) is now brought to self-hosted runners connected to GitHub Enterprise Server when installed via the Harden-Runner GitHub Action.

Run Harden-Runner on a GitHub Enterprise Server self-hosted ephemeral VM when baking the agent into the image is not possible.

```yaml
- uses: step-security/harden-runner@v2.22.1
  with:
    deploy-on-self-hosted-vm: true
    customer: your-stepsecurity-customer
    server-name: your-ghes-server
    egress-policy: audit
```

— Names exact inputs and agent version, no deeper mechanism given.

Harden Runner Releases: October 9, 2026…

---

03 **Runtime agent deployment on ephemeral self-hosted VMs** NEW

80

Adds `deploy-on-self-hosted-vm: true` input to deploy the agent directly on an ephemeral self-hosted Linux VM at workflow runtime, for cases where baking the agent into the VM image is not possible.

Run Harden-Runner on a GitHub Enterprise Server self-hosted ephemeral VM when baking the agent into the image is not possible.

```yaml
- uses: step-security/harden-runner@v2.22.1
  with:
    deploy-on-self-hosted-vm: true
    customer: your-stepsecurity-customer
    server-name: your-ghes-server
    egress-policy: audit
```

— Exact flag and runnable YAML snippet provided. Harden Runner Releases: October 9, 2026…

---

04 **macOS deny-list egress policy** IMPROVED

65

Extends the `denied-endpoints` input to macOS runners, enabling deny-list egress policy on macOS to block specific destinations while leaving all others reachable.

Block specific egress destinations on a macOS runner while leaving all other outbound traffic reachable.

```yaml
- uses: step-security/harden-runner@v2.22.1
  with:
    egress-policy: block
    denied-endpoints: 'example.com:443 suspicious-host.io:80'
```

— Names the flag and platform but limited mechanism detail.

Harden Runner Releases: October 9, 2026…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Email Security

Rspamd

Feedback ⓘ 4.2.2 NOTES ↗ RANK

Rapid spam filtering system.

> Rspamd 4.2.2 replaces Bitcoin-only wallet detection with multi-currency checksum verification and ships a cluster of breaking fixes to metadata_exporter validation, known_senders TTL enforcement, reply_sender_privacy, milter_headers scoping, and MIME part digests.

└──▷ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 x-spamd-bar scoped to skip_local/skip_authenticated

BREAKING

95

`x-spamd-bar` in `milter_headers` is now subject to `skip_local` and `skip_authenticated`, both `true` by default, so the header is suppressed for loopback, private-network, and authenticated-user messages. To keep the header for those messages, add `local_headers = ["x-spamd-bar"];` and `authenticated_headers = ["x-spamd-bar"];` to `local.d/milter_headers.conf`.

Restore x-spamd-bar for local and authenticated mail after milter_headers begins honouring skip_local/skip_authenticated for that header.

```hcl
local_headers = ["x-spamd-bar"];
authenticated_headers = ["x-spamd-bar"];
```

— Exact config keys, defaults and a runnable fix snippet product docs

02 MIME part and message digest recalculation

BREAKING

95

Digests for all non-base64-encoded parts in multipart messages change due to corrected RFC 2046 boundary handling, meaning fuzzy hashes learned before the upgrade no longer match and digest-based deduplication or joins across the upgrade boundary (e.g. a ClickHouse Digest column, Elasticsearch) will be incorrect. Use `rspamadm mime stat -F /path/to/message.eml` to inspect part digests as computed by 4.2.2.

Inspect part digests as computed by 4.2.2 before re-learning fuzzy samples, to confirm hashes have changed after the MIME boundary fix.

```
$ rspamadm mime stat -F /path/to/message.eml
```

— Names the command, affected systems and the RFC mechanism product docs

03 **Multi-currency crypto wallet detection** NEW DETECTION BREAKING

85

Replaces `rules/bitcoin.lua` with `rules/crypto.lua`, adding checksum-verified detection symbols `CRYPTO_ADDR_CHECK`, `LITECOIN_ADDR`, `DOGECOIN_ADDR`, `TRON_ADDR`, `XRP_ADDR`, `ZCASH_ADDR`, `CARDANO_ADDR`, `COSMOS_ADDR`, `STELLAR_ADDR`, `TON_ADDR`, plus format-only symbols `ETHEREUM_ADDR_MAYBE` and `MONERO_ADDR_MAYBE`. `LEAKED_PASSWORD_SCAM` (score 7.0) now fires for every checksum-verifiable currency instead of Bitcoin only; any custom Lua config loading `rules/bitcoin.lua` directly must be updated to reference `rules/crypto.lua`.

— Names every new symbol and the exact file rename but no migration command product docs

04 **known_senders TTL now enforced on upgrade** BREAKING

85

In `known_senders` zset mode (`use_bloom = false`), the first new sender recorded after upgrading to 4.2.2 removes every entry whose first-seen score is older than `max_ttl` (default `30d`), causing those senders to receive `UNKNOWN_SENDER` instead of `KNOWN_SENDER` until they send again. Set `max_ttl = 0;` in `local.d/known_senders.conf` to restore the pre-4.2.2 behaviour of never expiring entries.

Prevent known_senders entries from ever expiring after upgrading, preserving the pre-4.2.2 behaviour where senders are kept indefinitely.

```hcl
max_ttl = 0;
```

— Named config key, default value, and exact override snippet product docs

05 **Schema validation in metadata_exporter** BREAKING

75

`metadata_exporter` now validates configuration on load for the built-in backends `http`, `send_mail`, `redis_pubsub`, `redis_stream`, `redis_list`, and `json_raw_tcp`: rules with unknown keys or wrong-typed values are disabled rather than silently accepted, and the module itself is disabled if no valid rule remains. Look for `rule <name>: invalid configuration -> rule DISABLED` in logs and fix `local.d/metadata_exporter.conf`.

— Names backends and the log string but gives no fix command product docs

06 reply_sender_privacy honoured in replies and known_senders

70

BREAKING

`reply_sender_privacy`, `reply_sender_privacy_alg`, `reply_sender_privacy_prefix`, and `reply_sender_privacy_length` are now correctly read in both the `replies` and `known_senders` modules, which previously ignored them due to reading under the wrong config prefix. Because keys are now built from hashed addresses, reply sets and known-sender sets written by older versions are no longer found, so `INC_MAIL_KNOWN_LOCALLY` and `INC_MAIL_KNOWN_GLOBALLY` stop matching earlier correspondents until those senders send new mail.

— Names every affected config key but offers no remediation step product docs

└──▷ BREAKING ON UPGRADE

- ! `rules/bitcoin.lua` has been removed; any custom Lua configuration that loads it directly (e.g. a modified `rules/rspamd.lua`) will fail to load — change references to `rules/crypto.lua`.
- ! `metadata_exporter` rules that contain unknown keys or wrong-typed values for the built-in backends (`http`, `send_mail`, `redis_pubsub`, `redis_stream`, `redis_list`, `json_raw_tcp`) are now disabled on startup instead of running with the bad key ignored; check logs for `rule <name>: invalid configuration -> rule DISABLED` and fix `local.d/metadata_exporter.conf`.
- ! In `known_senders` zset mode (`use_bloom = false`), the first new sender recorded after upgrading to 4.2.2 will remove every entry whose first-seen score is older than `max_ttl` (default `30d`), causing those senders to receive `UNKNOWN_SENDER` instead of `KNOWN_SENDER` until they send again.
- ! If `reply_sender_privacy = true` was set in `local.d/replies.conf` or `local.d/known_senders.conf`, reply sets and known-sender sets written by older versions are no longer found because keys are now built from hashed addresses; `INC_MAIL_KNOWN_LOCALLY` and `INC_MAIL_KNOWN_GLOBALLY` stop matching earlier correspondents until users send new mail.
- ! `x-spamd-bar` in `milter_headers` is now suppressed for loopback, private-network, and authenticated-user messages (because `skip_local` and `skip_authenticated` are both `true` by default); Sieve rules or mail-client filters that relied on `x-spamd-bar` being present on those messages will stop seeing it.
- ! MIME part and message digests change for all non-base64-encoded parts in multipart messages; fuzzy hashes learned before the upgrade no longer match, and digest-based deduplication or joins across the upgrade boundary (e.g. in ClickHouse `Digest` column, Elasticsearch) will produce incorrect results.
- ! `LEAKED_PASSWORD_SCAM` (score 7.0) now fires for all checksum-verifiable currencies, not only Bitcoin; sites that tuned scores or composites assuming Bitcoin-only behaviour will see changed scoring for multi-currency scam messages.

4.2.2 NOTES ↗ LATEST STABLE

Adds richer metadata_exporter templating with custom variables, multipart email alerts, schema validation, and new Lua API helpers for header folding and SMTP negotiation.

└──▷ GET THIS VERSION

```
$ git clone --branch 4.2.2 https://github.com/rspamd/rspamd.git
# already have the repo? check out this version:
$ git checkout 4.2.2
```

Adds `mail_from`, `mail_to`, `helo`, `channel`, and `stream_key` fields plus selector expansion and custom variables to `metadata_exporter` rules.

Adds `email_parts` to `metadata_exporter` to compose multipart email alerts with plain and HTML parts grouped into `multipart/alternative`, with automatic RFC 2047 encoding for non-ASCII headers.

Validates `metadata_exporter` rules against a schema so that `configtest` fails on an invalid rule.

Adds `util.fold_header_unstructured` to the Lua API for folding unstructured headers per RFC 5322.

Adds an `smtp` newline normalization mode to `rspamd_text` in the Lua API.

+3 more

└──▷ BREAKING ON UPGRADE

! The CRLF before a boundary delimiter no longer counts as part of the preceding MIME part — message digests and non-base64 part digests change for all multipart messages, so anything keyed on those values will not match values computed by versions before 4.2.2.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Cloudmersive

Feedback ⓘ

Cloudmersive Cloud Fraud AI API v1-7-0 Released (2026-10-10) BLOG ↗ RANK

Cloudmersive is a developer platform offering cloud APIs for file scanning, malware detection, document conversion, image processing, and other validation and data-processing tasks.

> Cloudmersive's Fraud AI API gained a setting to allow or block expired documents during fraud detection checks.

└──▷ **WHAT SHIPPED · 1 FEATURE** most completely described first ? what's the number?

01 **Expired-document handling in Fraud AI checks** NEW 30

The Cloud Fraud AI API adds settings to allow or block expired documents during fraud detection checks.

— Names the setting's purpose but no field name or mechanism

Cloudmersive Cloud Fraud AI API v1-7-0 …

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

syzkaller

Feedback ⓘ changes since 2026-09-24 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller gains batch task execution for syz-aflow, multi-API-key Gemini support, and new Linux syscall descriptions for Landlock and SCTP.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-09-24
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-24
```

Adds batch task execution to `tools/syz-aflow`, allowing multiple `.json` task files from a directory to be run in parallel with per-task log isolation, records executed seeds as trajectory span artifacts, and collects results and executed programs into a `corpus.db`.

Adds support for multiple API keys in `pkg/aflow/backend/gemini`, enabling load distribution across Gemini API keys.

Adds `sys/linux` descriptions for the Landlock `no_new_privs` flag (with a corresponding test in `sys/linux/test`), missing SCTP socket options, and `rnull` configfs device creation, and improves Rust binder descriptions in `sys/linux`.

Adds auto-detection of CPU vendor and nested virtualization for boot tests in `syz-manager` and `pkg/vminfo`, and adds verification of `required_test_capabilities` in `ModeRunTests` within `syz-manager` and `pkg/mgrconfig`.

Adds support for batched searches and per-file counts in `pkg/aflow/tool/grepper`.

+4 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue