

Telemetry sources

System logs (text)

auth_log
syslog
web_log

eBPF kernel events

process_events
network_events
file_events

Observation

Dict over channels
modes:
Text | Hybrid

o_t

Agent

streaming continual learner
no replay buffer
predict-then-update

a_t

Action (Dict)

discrete \in
{pass, alert,
throttle,
block_source,
unblock,
isolate}
 \times risk $\in [0, 10]$

Defense state

blocklist
throttle list
isolation mode

Reward

asymmetric per-action term
+ ongoing consequences

r_t

consequences

causally suppresses future observations

Continuous, non-episodic stream (terminated = False)