Stopping web attacks inline, on CSIC 2010

61,000 HTTP requests to one application, 25,065 of them attacks. The engine learns from normal requests only and never sees an attack during learning.

91.69% of attacking sources blocked by their third request, with 0 of 6,000 genuine sources wrongly blocked

RESULTS

REQUESTS SEEN FROM A SOURCE	ATTACKING SOURCES BLOCKED	GENUINE WRONGLY BLOCKED
2 requests	73.38%	0 of 9,000
3 requests	**91.69%**	**0 of 6,000**
5 requests	99.68%	0 of 3,600
10 requests	100%	0 of 1,800

Bars run 0 to 100%.

HOW IT WAS MEASURED

Unit	The source: a run of consecutive requests from one client
Atom	The raw bytes of each request. No parsing and no signatures
Learning	Normal requests only; no attack is labelled or seen
Level	Set on normal traffic alone

CONDITIONS

Single-request ranking AUC	0.9221
Decision cost	about 0.09 ms per request, about 11,000 requests a second on one core
Peak memory	29 MB
GPU	none

THE PASS MARK, FIXED BEFORE THE RUN

> Set before the run: 90% or more of attacking sources blocked with 1% or fewer genuine sources wrongly blocked. Verdict: PASS.

LIMITS OF THIS RESULT

CSIC 2010 carries no client addresses, so a source is a run of consecutive requests of one class; a real attacker may mix in normal requests. It is one application and generated traffic from 2010. Judging by source cannot act on an attack that uses each address once.

WHERE IT IS RECORDED

simplex/validate/csic2010/CSIC_FINDINGS.md; showcase run 2026-09-19. The data behind this sheet is published beside it at senuaai.com/evidence/datasheets/.