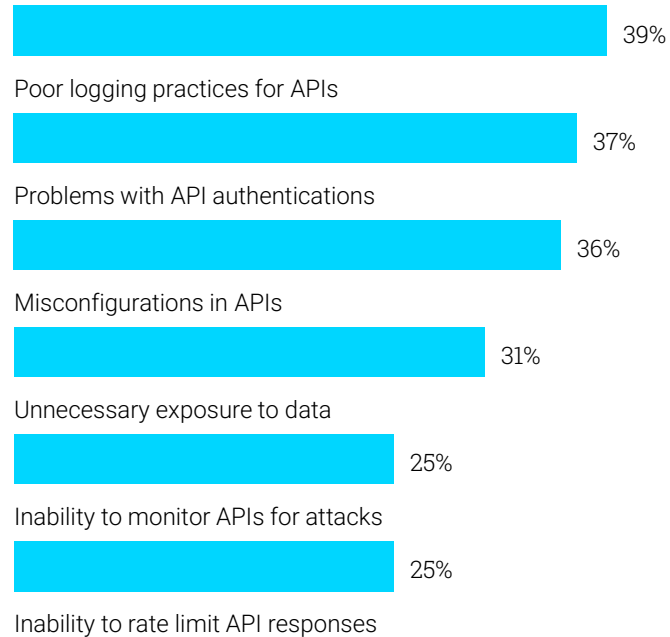


API Security requires Machine Learning, because it is a superhuman problem to solve



15,564

Average number of Production Enterprise APIs

76%


of organizations experienced a security breach
in the past year

37 days

27 days for discovery

10 days for remediation

per incident



Whitepaper

The 2022 API Security Trends Report

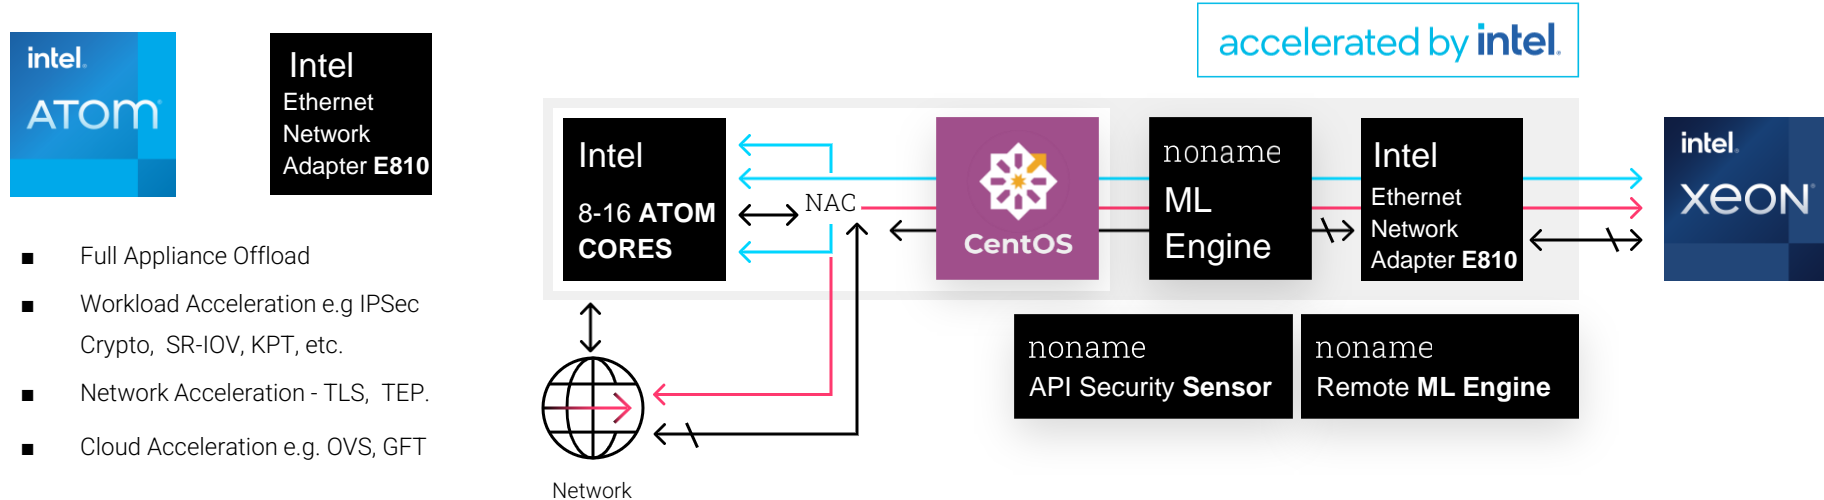
[Learn more](#) →

Intel® NetSec Accelerator Reference Design

with 100GbE Intel® Ethernet Network Adapter E810 Edge Telco Grade network interface & low latency API Security powered by Noname Security Machine Learning

Family of Intel SoC-based Reference designs intended to accelerate networking and security workload

Use cases: IPSec/SSL/Http offload, SASE, FW | **AI Machine Learning integration with:**



A complete autonomous server with full Orchestration and Management on a PCIe add-in card

The screenshot shows the Noname Security dashboard. The top navigation bar includes the Noname logo, menu items for Security, Inventory, Reports, and Testing, and a button for API Specs. A search bar for engines is located at the top right of the main content area. The left sidebar contains navigation options: Datatypes, Authentications, Policies, and Integrations. The main content area is titled 'Engines' and features two buttons: 'Add Engine' and 'Download Package'. Below these is a table listing the installed engines.

Engine ⓘ	Hosted	Status ▼
default-engine Integrated	-	● Online
IntelNetSecAcceleratorE810PhantomLake+Noname Remote	-	● Online

```

On-line CPU(s) list: 0-15
Thread(s) per core: 1
Core(s) per socket: 16
Socket(s): 1
NUMA node(s): 1
Vendor ID: GenuineIntel
BIOS Vendor ID: Intel(R) Corporation
CPU family: 6
Model: 134
Model name: Intel Atom(R) P5742 processor
BIOS Model name: Intel Atom(R) P5742 processor
Stepping: 7
CPU MHz: 1077.306
CPU max MHz: 2200.0000
CPU min MHz: 800.0000
BogoMIPS: 4400.00
Virtualization: VT-x
L1d cache: 32K
L1i cache: 32K
L2 cache: 4608K
L3 cache: 12800K
NUMA node0 CPU(s): 0-15
Flags: fpu vme de pse tsc msr pae mce cx8 apic
n pbs bts rep_good nopl xtopology nonstop_tsc cpuid aperfmp
dline_timer aes xsave rdrand lahf_lm 3dnowprefetch cpuid fau
rms cqm rdt_a rdseed smap clflushopt clwb intel_pt sha_ni xs
id cldemote movdiri movdir64b md_clear flush_lld arch_capabi
[root@ptl ~]# cat /etc/os-release
NAME="CentOS Linux"
VERSION="8"
ID="centos"
ID_LIKE="rhel fedora"
VERSION_ID="8"
PLATFORM_ID="platform:el8"
PRETTY_NAME="CentOS Linux 8"
ANSI_COLOR="0;31"
CPE_NAME="cpe:/o:centos:centos:8"
HOME_URL="https://centos.org/"
BUG_REPORT_URL="https://bugs.centos.org/"
CENTOS_MANTISBT_PROJECT="CentOS-8"
CENTOS_MANTISBT_PROJECT_VERSION="8"
    
```

CONTAINER ID	NAME	CPU %	MEM USAGE / LIMIT	MEM %
f3fb7cf4ac64	noname_router.1.pacr5xv2jwoqla5vcnd651te6	0.05%	173.7MiB / 31.05GiB	0.55%
692b7a13debd	noname_engine.1.uq4rayf5typ7zdd8gxml2uada	0.44%	238.9MiB / 16GiB	1.46%
d8f8ec513528	noname_nats.1.b4uvlk9g8lniwjm9aosrlzsv0	0.02%	18.01MiB / 31.05GiB	0.06%
980c0a6f7b1d	noname_nginx.1.s2cxji2qsz9phqlepszowz4cp	0.01%	29.17MiB / 31.05GiB	0.09%
e720bbc4c165	noname_syslog-ng.1.ly54v7dbg9ejxjhrbdqgcw9nx	0.00%	3.684MiB / 31.05GiB	0.01%


```

[root@ptl ~]# docker ps
CONTAINER ID   IMAGE                                COMMAND                  CREATED        STATUS
f3fb7cf4ac64   localhost/router:v3.9.2             "/app/entrypoint.sh ..." 5 hours ago   Up 5 hour
692b7a13debd   localhost/engine:v3.9.2             "/project/entrypoint..." 5 hours ago   Up 5 hour
d8f8ec513528   localhost/nats:2.8.4                 "/nats-server --conf..." 5 hours ago   Up 5 hour
980c0a6f7b1d   localhost/nginx:nms-1.23.3          "/docker-entrypoint..." 5 hours ago   Up 5 hour
e720bbc4c165   localhost/syslog-ng:balabit-3.38.1  "/usr/sbin/syslog-ng..." 5 hours ago   Up 5 hour
    
```

12 hours of JPEG data processed by noname machine learning on a 100 Gbps network card


The average CPU usage of the 16 atom cores was less than 10% while doing machine learning on network traffic and the temperature of the card remained cool at less than 40 degrees C, the value of doing machine learning on the network card is there is no transfer of API logs nor latency added on the API server resulting in improved performance and security and privacy

12 hours results

- Average response time = **3.12 ms**
- Fastest response time = **503 microseconds**
- Throughput = **43 Million API requests**
- At a rate of = **996 per second**
- Data sent = **14 GB**
- Data received = **1.5 TB**

```
CONTAINER ID   NAME                                     CPU %     MEM USAGE / LIMIT   MEM %     NET I/O     BLOCK I/O   PIDS
b7e92126174a   noname_engine.1.g3jhr8yjeo9w2cljk8alpkf6x  8.07%    2.889GiB / 16GiB    18.05%    0B / 0B     2.92MB / 1.02GB  31
8f34a6330949   noname_syslog-ng.1.5bsjms04pl0eusj0f2lh1081  0.00%    3.992MiB / 31.05GiB  0.01%    0B / 0B     401kB / 0B      3
8246b5f1ced0   noname_router.1.oin01jn2vjxqzlj6kovgisyhc3  0.07%    189.1MiB / 31.05GiB  0.59%    0B / 0B     340kB / 0B      33
ca34d28dc356   noname_nginx.1.12jq1uud9ft4kfwicedteaeav  0.01%    31.39MiB / 31.05GiB  0.10%    0B / 0B     0B / 20.5kB     20
5593a6147725   noname_nats.1.w813s0pizu44k2kqdlcoo76ha    0.02%    15.08MiB / 31.05GiB  0.05%    0B / 0B     0B / 0B         19
C
[root@pt1 ~]# uptime
 04:08:07 up 13 days, 6:02, 2 users, load average: 0.37, 0.32, 0.28
[root@pt1 ~]# 20 minutes of 1000 RPS get image jpeg from httpbin api server
bash: 20: command not found...
[root@pt1 ~]# sensors
acpitz-virtual-0
Adapter: Virtual device
temp1:          +26.8°C (crit = +100.0°C)

coretemp-isa-0000
Adapter: ISA adapter
Package id 0:  +40.0°C (high = +80.0°C, crit = +100.0°C)
Core 8:        +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 9:        +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 10:       +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 11:       +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 12:       +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 13:       +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 14:       +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 15:       +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 16:       +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 17:       +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 18:       +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 19:       +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 20:       +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 21:       +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 22:       +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 23:       +34.0°C (high = +80.0°C, crit = +100.0°C)
```

 Terminal - h-r750-01 (zir11) - zvlh32 -- NoName

```
Disconnect Connected - root@h-r750-01: /usr/local/mesh_services/meshagent
CONTAINER ID   NAME                                     CPU %     MEM USAGE / LIMIT   MEM %     NET I/O     BLOCK I/O   PIDS
2568484e3141   httpbin.alpine  78.31%    65.29MiB / 503.2GiB  0.01%    102GB / 146GB  8.19kB / 0B      2
```

1 hour of credit card data processed by noname machine learning on a 100 Gbps network card

The average CPU usage of the 16 atom cores was less than 25% while doing machine learning on network traffic and the temperature of the card remained cool at less than 40 degrees C, the value of doing machine learning on the network card is there is no transfer of API logs nor latency added on the API server resulting in improved performance and security and privacy

1 hour results

- Average response time = **8.63 ms**
- Fastest response time = **708 microseconds**
- Throughput = **3.56 Million API requests**
- At a rate of = **989 per second**

```
CONTAINER ID   NAME                                     CPU %     MEM USAGE / LIMIT   MEM %     NET I/O     BLOCK I/O   PIDS
b7e92126174a noname_engine.1.g3jhr8yjeso9w2c1jk8aipkfkx 24.22%    3.617GiB / 16GiB    22.61%    0B / 0B     1.92MB / 1.02GB 31
8244e6330949 noname_syslog-ng.1.5bsjms94p10eusj0f2lh1081 0.00%     3.459MiB / 31.05GiB 0.01%     0B / 0B     0B / 0B      3
8246b5f1ced0 noname_router.1.o1n0l1ja2vjzqlj6kcvogsyhc3 0.05%     187.8MiB / 31.05GiB 0.59%     0B / 0B     328kB / 0B    33
ca34d28dc356 noname_nginx.1.12jqluvd9ft4kfzwicdeaeav    0.01%     31.02MiB / 31.05GiB 0.10%     0B / 0B     0B / 20.5kB   20
5593a6147725 noname_nats.1.w813s0pizu44k2kqdlcoo76ha    0.03%     17.7MiB / 31.05GiB 0.06%     0B / 0B     0B / 0B      18

[root@ptl ~]# sensors
acpitz-virtual-0
Adapter: Virtual device
temp1:      +26.8°C (crit = +100.0°C)

coretemp-isa-0000
Adapter: ISA adapter
Package id 0: +40.0°C (high = +80.0°C, crit = +100.0°C)
Core 8:      +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 9:      +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 10:     +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 11:     +33.0°C (high = +80.0°C, crit = +100.0°C)
Core 12:     +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 13:     +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 14:     +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 15:     +32.0°C (high = +80.0°C, crit = +100.0°C)
Core 16:     +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 17:     +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 18:     +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 19:     +36.0°C (high = +80.0°C, crit = +100.0°C)
Core 20:     +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 21:     +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 22:     +34.0°C (high = +80.0°C, crit = +100.0°C)
Core 23:     +34.0°C (high = +80.0°C, crit = +100.0°C)
```



Name	RPS	Total Requests	Last traffic receive	Status
IntelNetSecAcceleratorE810NonameRemoteEngine	991.70	1.65M	57 seconds...	Online



Category Leader

- Largest API security company
- \$200+ million in funding from leading investors
- Award-winning comprehensive platform, from code to run-time



Global Presence

- 300+ employees
- NA, EMEA, APJ & IL locations
- Follow-the-sun 24-7 support
- SOC2, STAR Level 1, PCI-DSS Compliant



Strong Partner Ecosystem

- Alliances with all major cloud providers and key API platforms
- Global GTM partnerships
- 75+ certified integrations with network security, CI/CD tools, SEIM, and workflow platforms



Industry-Leading Customers

- Working with over 20% of Fortune 500
- 3 of 5 largest US telcos
- 2 of 5 largest pharmaceuticals
- Major global financial, insurance, consumer & entertainment leaders





Noname Machine Learning Accelerated by Intel Confidential Computing with Intel® Trust Authority Attestation for Zero Trust AI Augmented API Security on Trustworthy Cloud Computing Infrastructure at global scale with telco grade performance.

Implement advanced AI augmented API Security innovations with Intel Confidential Computing trustworthy hybrid cloud infrastructure for regulatory compliance, data security, breach prevention, and data isolation. Confidential computing is already a government mandate in some countries for patient data compliance purposes in the Healthcare industry to protect ePHI.

Noname Security solves the superhuman problem of API Security in the cloud and on premises and we do it with the help of Intel Project Amber and the Intel Trust Authority service API to attest that the cloud or on-premises server hardware is equipped with genuine Intel® Trust Domain Extensions (Intel® TDX) confidential computing encryption enabled processors.

Also consider the possibilities of leveraging machine learning embedded innovations on a 100 Gbps system on a chip Intel NetSec Accelerator Reference Design Ethernet card.

noname

Thank You

For more information visit
nonamesecurity.co

m

