Building Security at Scale

Stanford Computer Forum 2014

Alex Stamos

Yahoo!

Who am I?

- 36 days as CISO of Yahoo
- Founder of Artemis Internet
- Co-Founder of iSEC Partners
- @stake, Loudcloud
- Cal BS EECS ‘01, worked on Patterson team

Take-Aways from Today

1. Internet-scale companies have unique economic security drivers
2. The security industry does not serve us well
3. Most academic research does not help
4. There is a huge opportunity for both academia and industry to work with us
5. Our problems will be everybody’s problems soon

When you think of an industry that is subject to online attacks, what first comes to mind?

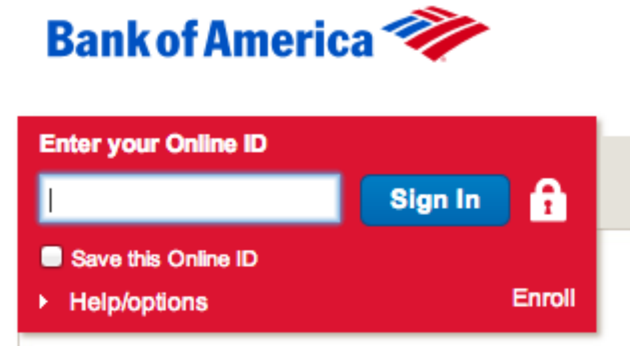

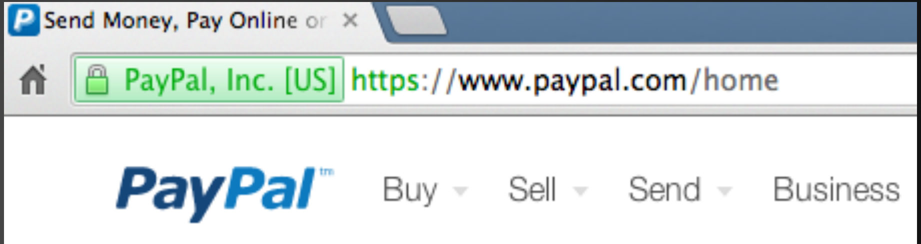
Send Money, Pay Online or
PayPal, Inc. [US] https://www.paypal.com/home
PayPal
Buy
Sell
Send
Business

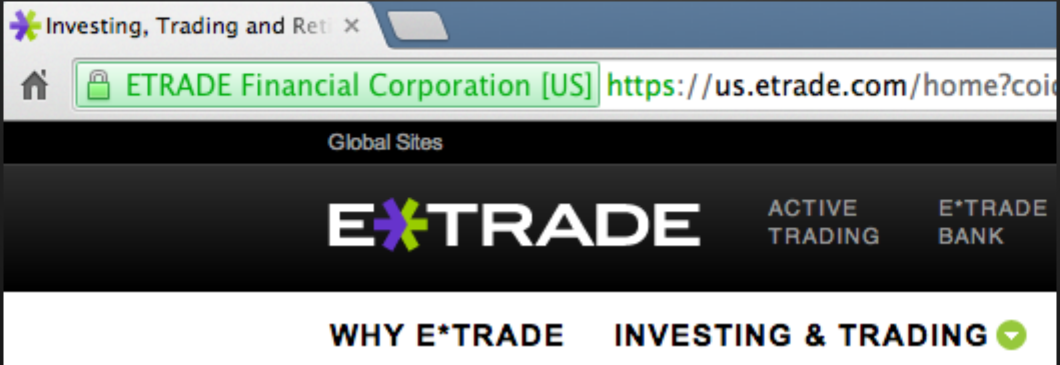
Investing, Trading and Ret
ETRADE Financial Corporation [US] https://us.etrade.com/home?coi
Global Sites
E*TRADE
ACTIVE TRADING
E*TRADE BANK
WHY E*TRADE
INVESTING & TRADING

How are these firms related?

Millions of customers

pay dozens to *hundreds* of dollars

visit *rarely*

and

have meat-space identities

How about for the Web Scale Companies?

Billions of customers

pay *nothing* (but click on ads)

visit *often*

and

have *no* link to real-life

	Big Banks	*Online Payments*	*Web Scale*
# of Customers	x 10^7	x 10^8	x 10^9
# of Concurrent Users	x 10^4	x 10^5	x 10^8
# of FE Servers	x 10^2	x 10^3	x 10^4
# of Total Servers	x 10^4	x 10^4	x 10^5
Customer Value	$100’s	$10’s	$.01s
Cust Stickiness	High	Medium	Low-Medium
Meat-Space Identity	Strong	Moderate	Weak
Post-Facto Action?	Yes	Yes	Rarely

Two totally different problems:

Banks are protecting real customers from attack

Web companies are trying to figure out which users are assets and which are liabilities

Things people try to sell us

Things people try to sell us: Smart Firewalls!

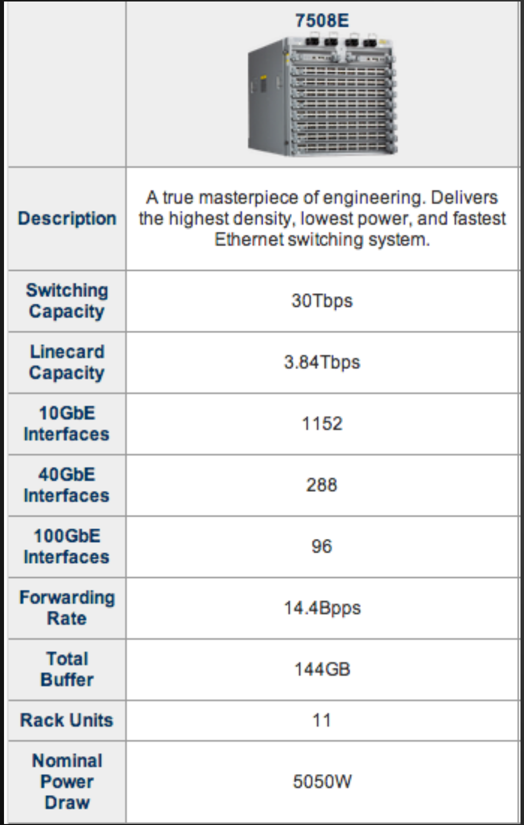

	7508E
Description	A true masterpiece of engineering. Delivers the highest density, lowest power, and fastest Ethernet switching system.
Switching Capacity	30Tbps
Linecard Capacity	3.84Tbps
10GbE Interfaces	1152
40GbE Interfaces	288
100GbE Interfaces	96
Forwarding Rate	14.4Bpps
Total Buffer	144GB
Rack Units	11
Nominal Power Draw	5050W

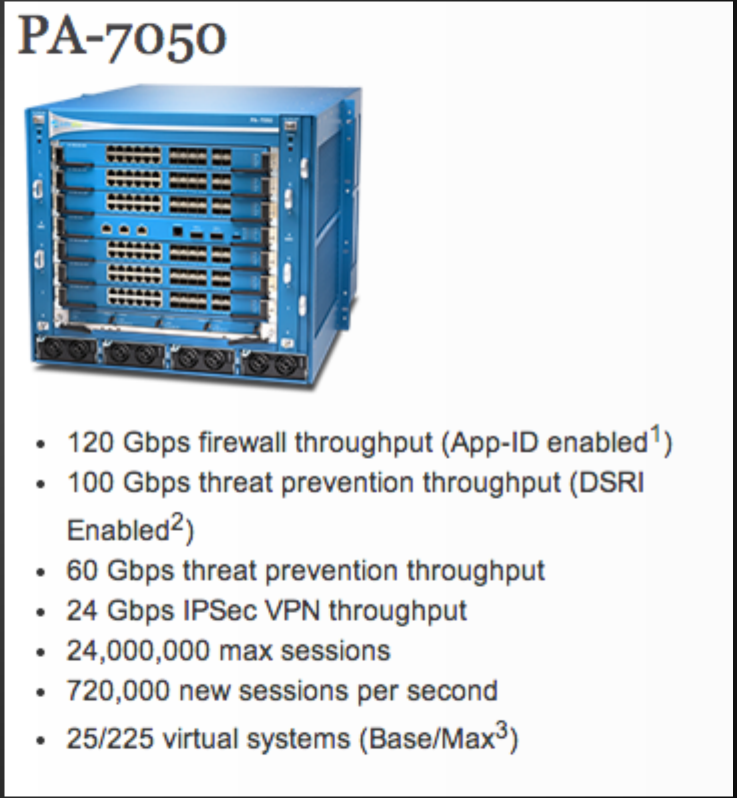

PA-7050

- 120 Gbps firewall throughput (App-ID enabled[1])
- 100 Gbps threat prevention throughput (DSRI Enabled[2])
- 60 Gbps threat prevention throughput
- 24 Gbps IPSec VPN throughput
- 24,000,000 max sessions
- 720,000 new sessions per second
- 25/225 virtual systems (Base/Max[3])

5kW

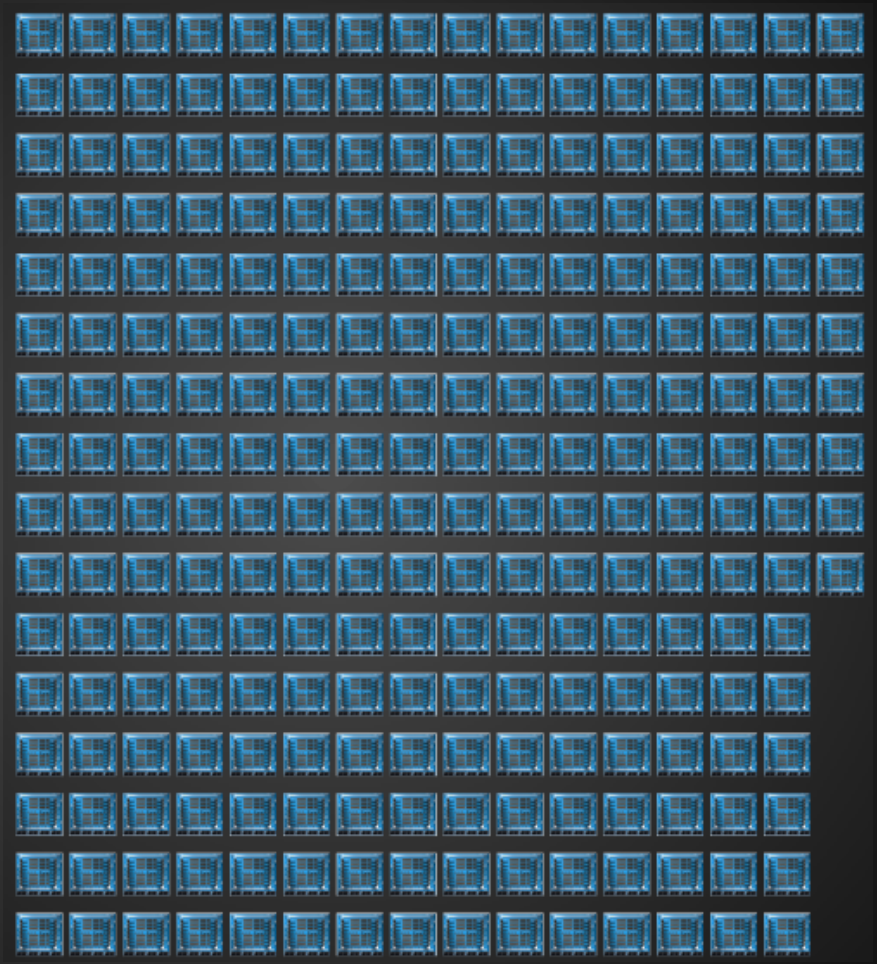

600kW

Things people try to sell us: DB Backed SIEM

"Just dump your data into Oracle and put it on a SAN"

EMC^2

1000GB 7200RPM SATA II	$1,450

1PB=$1.4M

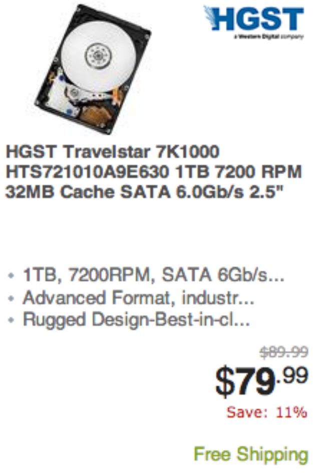

HGST
a Western Digital company

HGST Travelstar 7K1000 HTS721010A9E630 1TB 7200 RPM 32MB Cache SATA 6.0Gb/s 2.5"

- 1TB, 7200RPM, SATA 6Gb/s...
- Advanced Format, industr...
- Rugged Design-Best-in-cl...

~~$89.99~~

$79.99

Save: 11%

Free Shipping

1PB=$80K

Things people try to sell us: IDS Appliances

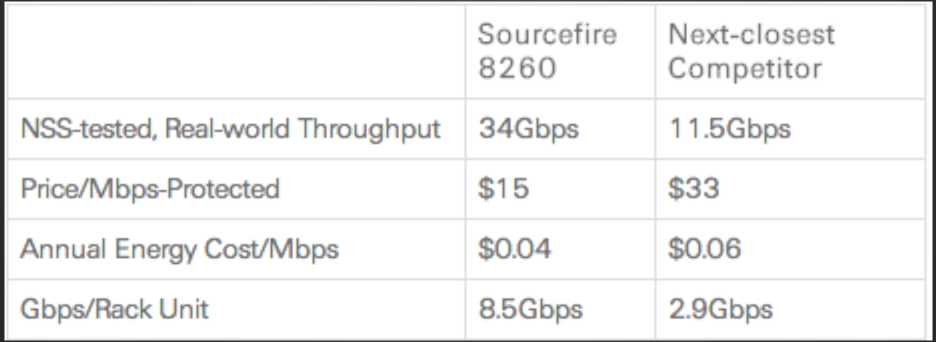

	Sourcefire 8260	Next-closest Competitor
NSS-tested, Real-world Throughput	34Gbps	11.5Gbps
Price/Mbps-Protected	$15	$33
Annual Energy Cost/Mbps	$0.04	$0.06
Gbps/Rack Unit	8.5Gbps	2.9Gbps

300Gbps x $15/Mbps = $4.5M

Things people try to sell us: Reputation Services

“Call our web service with the data and we’ll return a result in only 2000ms.”

In an ad-supported business, latency is death.

Aren’t we a special case?

Not really…

- Big data means that power efficiency is becoming a competitive advantage for many
  - Finance
  - Biotech
  - Logistics and Operations
- Latency is also more important than ever
  - See “Flash Boys” by Michael Lewis

Where security needs to go

Collapse the perimeter

Security services need to be as close as possible to the data you are protecting:

- Anomaly/Intrusion Detection
- Data Encryption
- AAA
- Network access control

Only sell software. Pizza boxes are great for pizza.

False Positives are Death

.01% False Positive Rate x 800M MAU =

80,000 alerts

- Alerting isn’t my problem
- The response funnel needs to narrow quickly

Latency is Death

Security needs to move towards asynchronous reactions

DRM world provides good examples

Better Mousetraps

Freemium Key Management and App Auth

Dual-auth TLS is the future of app auth

- Conceptually simple
- Open-source foundation
- Decentralized failure modes
- X.509 reasonably flexible

Why isn’t there a MySQL for Auth?

Bug Bounty with Automated Verification

Bug bounties create huge problems for companies

Why can’t the reporter upload a selenium script that verifies the issue?

Solving this would open SME market

Reputation Services that Work

Industry is moving to “slow auth”

What I want is:

- Open-Source
- To benefit from other sensors with privacy
- Realistic geo-based tracking
- Accuracy with IPv6

Hadoop Based SIEM

Proprietary distributed file systems will die

- Let me figure out how I store my data

Give me:

- Fast scrubbing/tokenization
- Natural language search
- Useful visualization
- Pre-defined but tunable anomaly models

ARM Based Secure Systems

ARM is going to take over the datacenter

If you could go back and build the x86 datacenter, what would you do?

- Lightweight containers
- Aggressive anti-exploit
- Trusted, diskless boot

Thank you!

alex@stamos.org

stamos@yahoo-inc.com