Scanning Darkly

DERBYCON

Hello Derbycon

HD Moore

- Metasploit founder and chief architect
- Chief research officer for Rapid7
- Head of Rapid7 Labs

Twitter: @hdmoore

Email: hdm@rapid7.com

RAPID7

zo*hB3LoleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!Tr0ub4dor&rzo*hB3LoleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

Derbycon 1.0

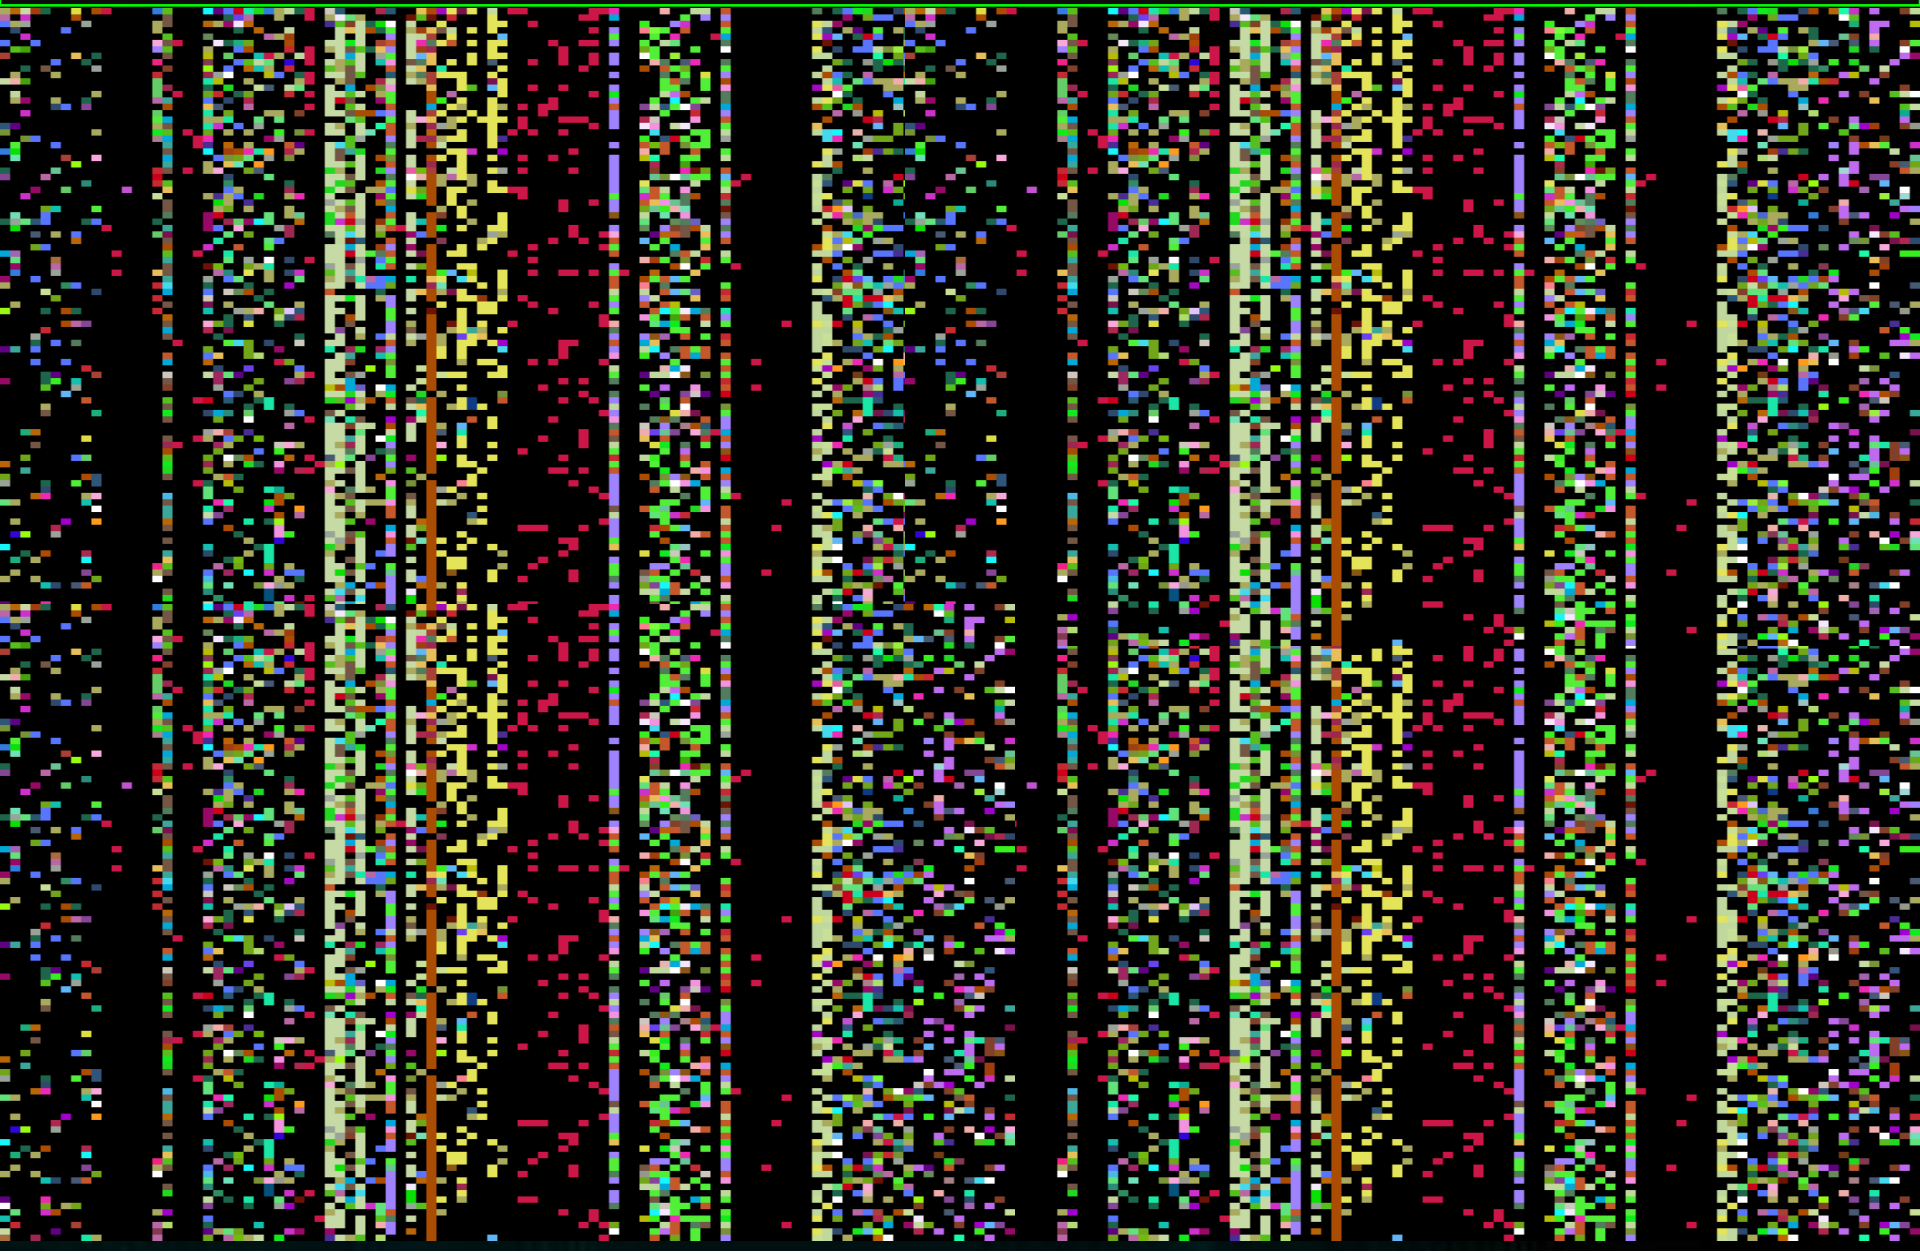

zo*h83LoleLAuNodTwuNz7$%i9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!Tr0ub4dor&rzo*h83LoleLAuNodTwuNz7$%i9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

Derbycon 1.0

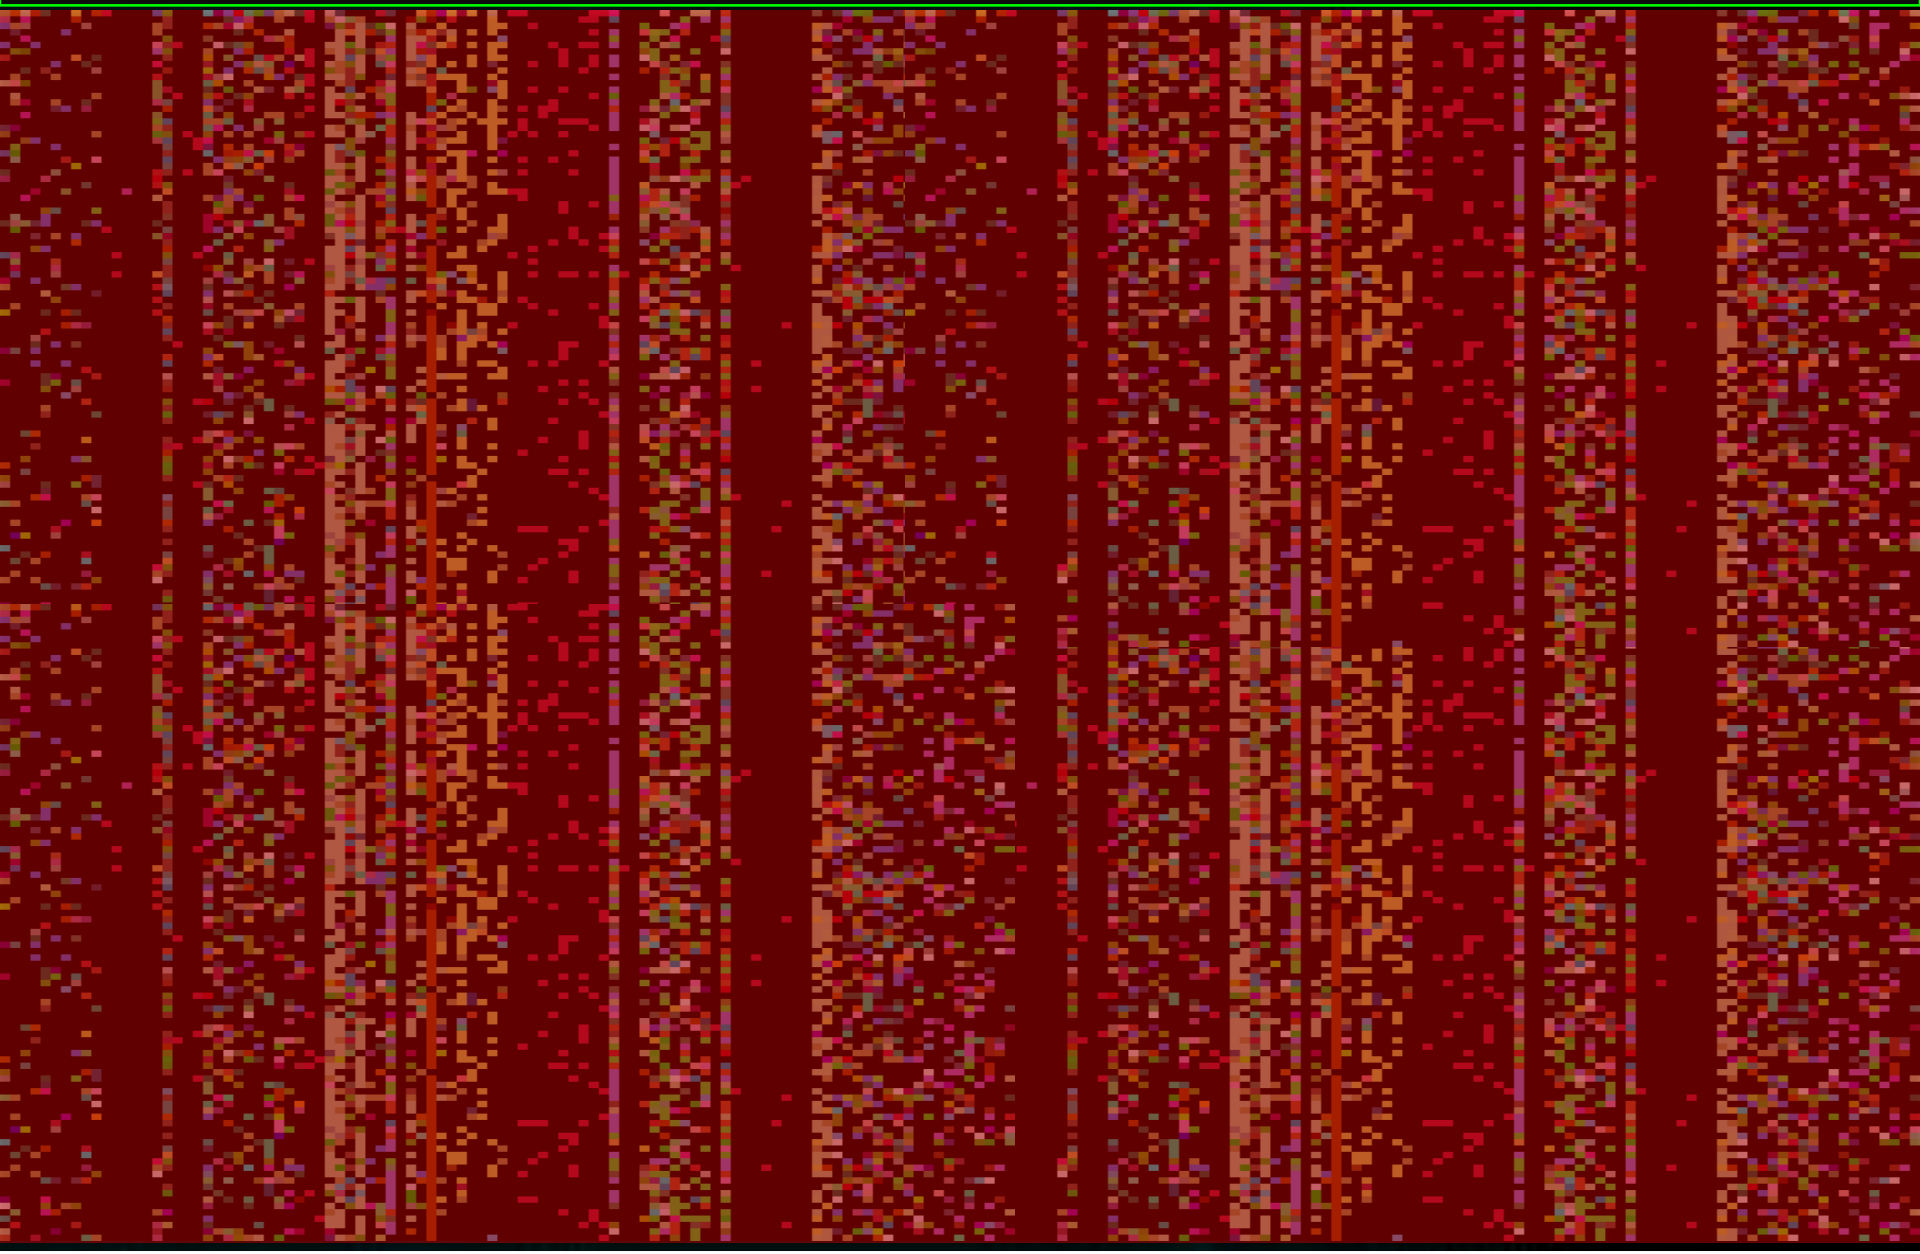

zo*hB3LoleLAuNodTwuNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LoleLAuNodTwuNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

Derbycon 2.0

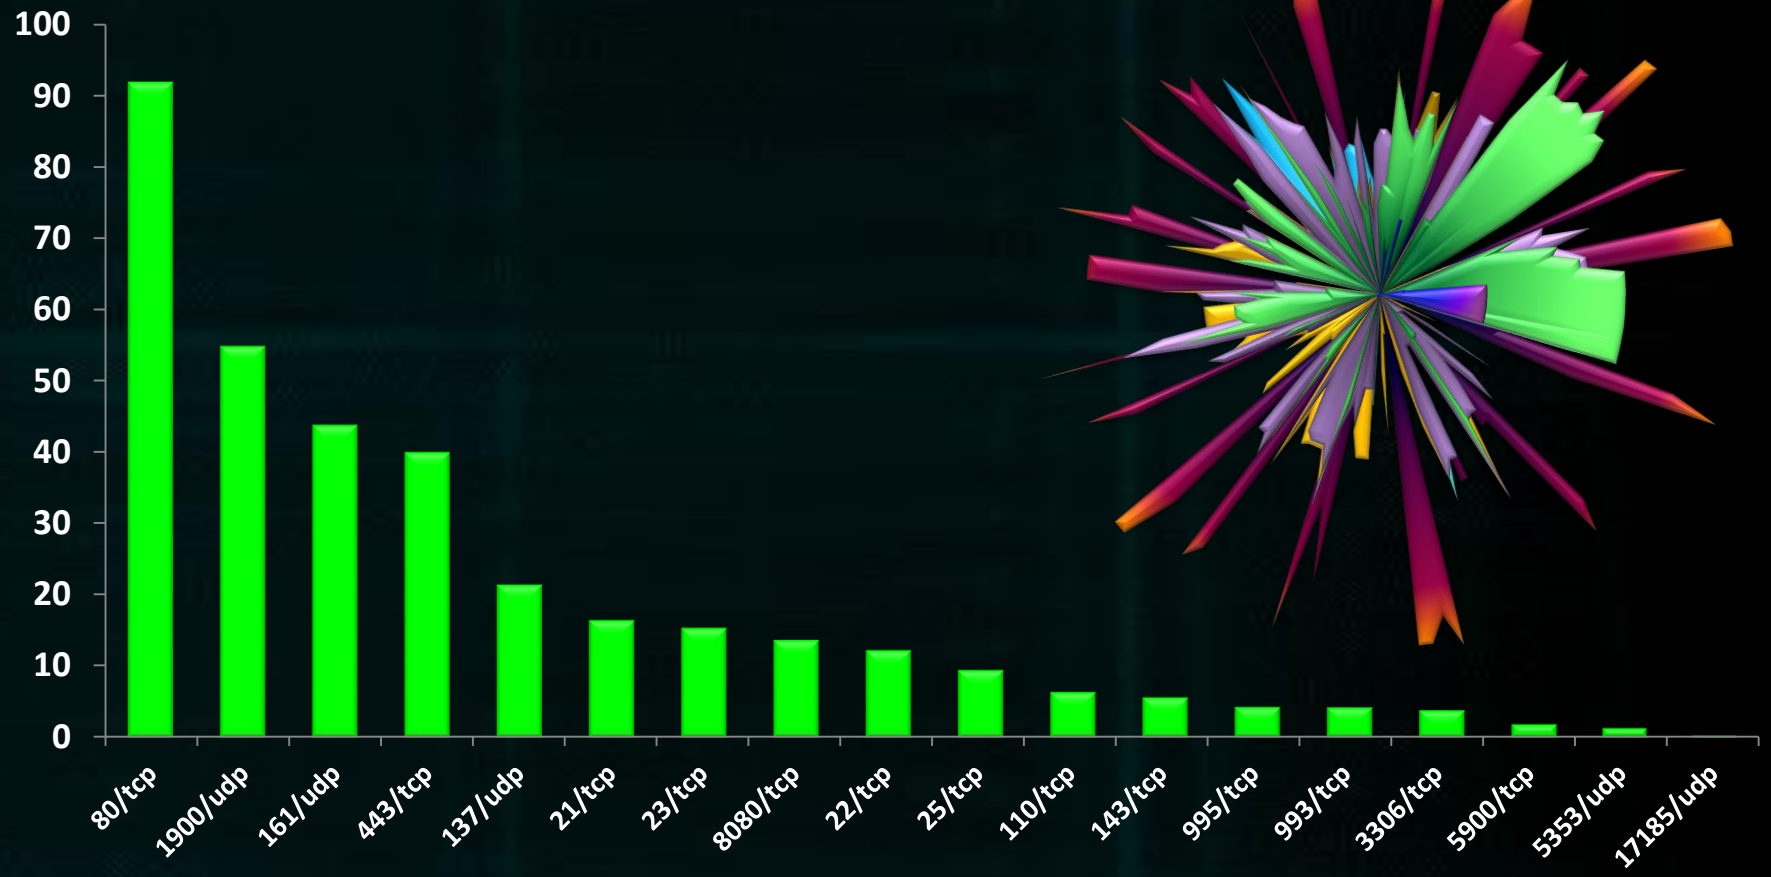

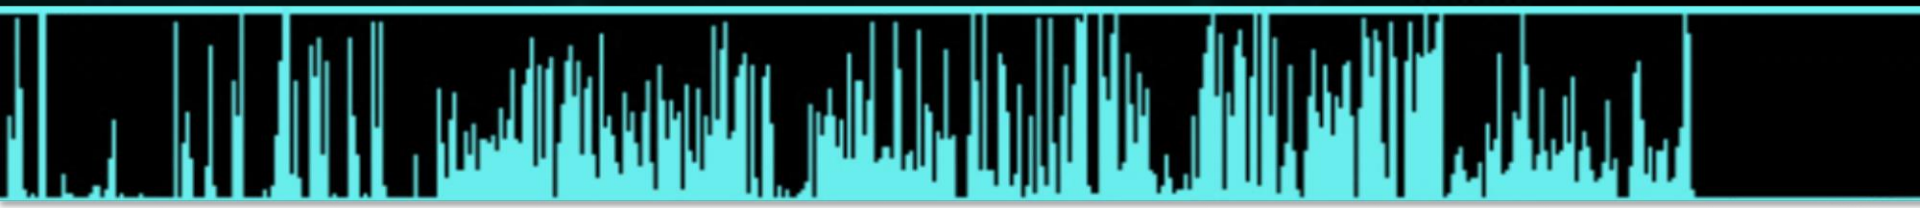

zo*hB3LoleLAuNodTwNz7$%:i5iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LoleLAuNodTwNz7$%:i5iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

Derbycon 2.0

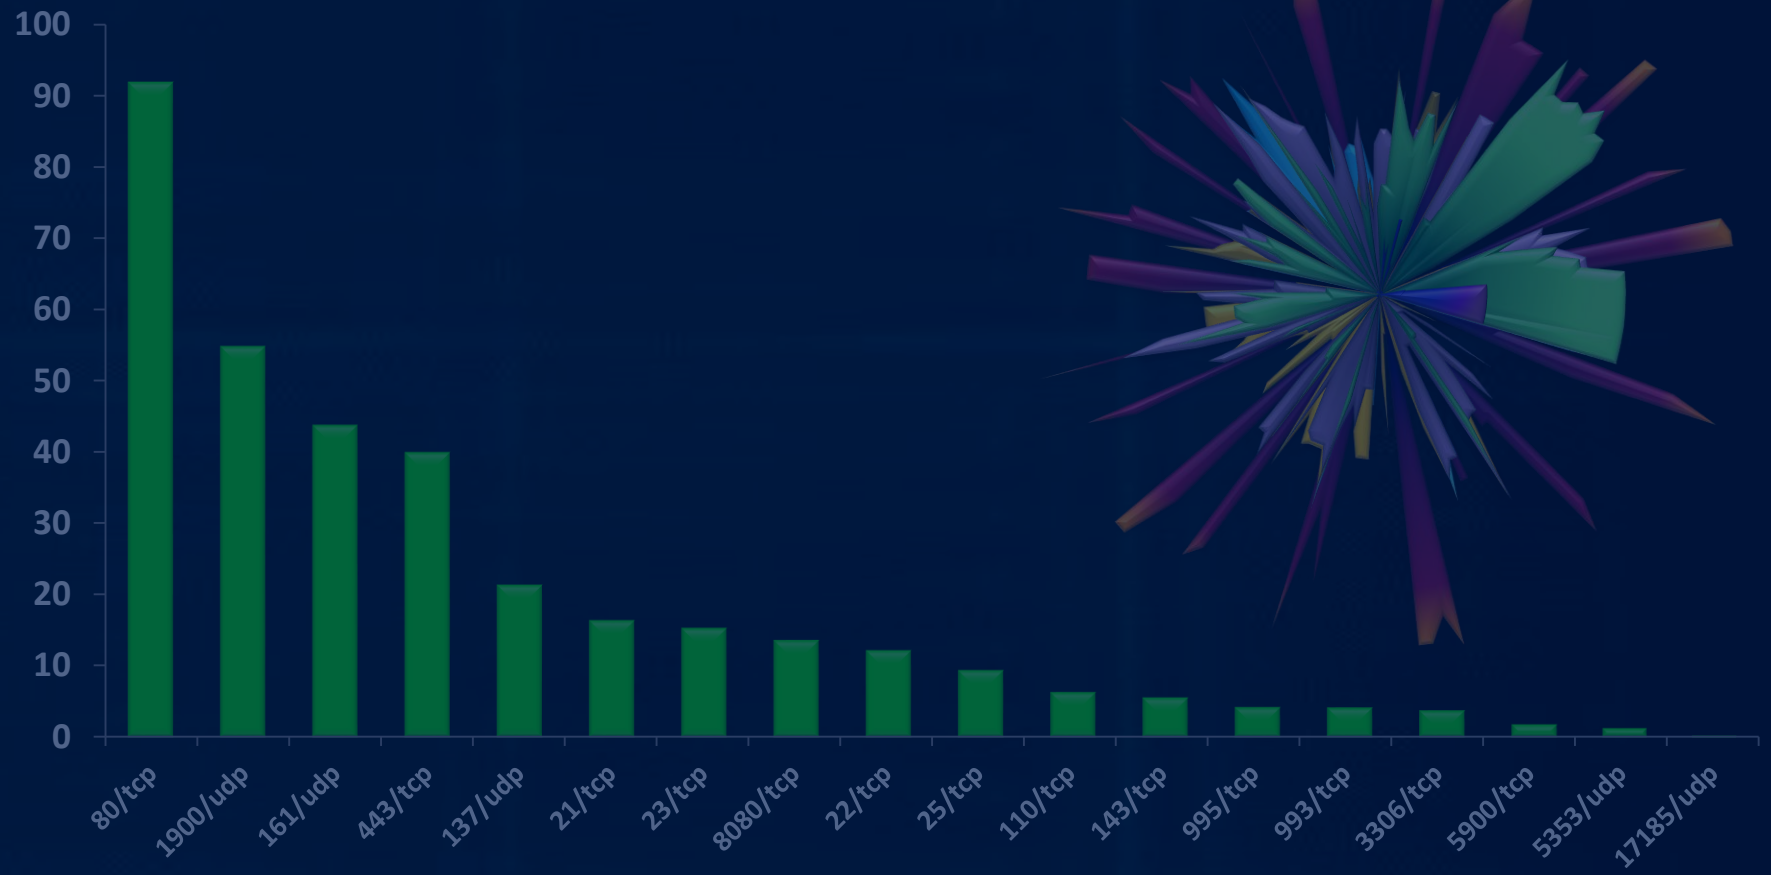

zo*hB3LaleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LaleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

Mass scanning is starting to mature

- Major improvements to scanning tools
- Numerous large-scale scanning efforts
- Scary and not-so-scary precedents

zo*hB3LoleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LoleLAuNodTwNz7$/:9iUDtsDt1spBu9DtQPEb2comBpua$NTxBT*kQ9*!

ZMap

U. Michigan team released Zmap

- Send a single probe across IPv4 in 45 minutes
- Detailed research paper with examples
- Development continues at GitHub
- Epic forge-socket support
- http://zmap.io


```
$ zmap -p 80 -o results.txt
```

ZMap: Data Collection

Over 110 internet-wide SSL scans in 12 mos

- Created a detailed view of the SSL ecosystem
- Realtime monitoring of Sandy outages
- Obtained 43 million unique certs

zo*hB3LoleLAuNodTwNz7$/:5iUDtsDt1spBu9DtQPEb2comBpuo$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LoleLAuNodTwNz7$/:5iUDtsDt1spBu9DtQPEb2comBpuo$NTxBT*kQ9*!

MASSCAN

Errata Security released Masscan

- Scan all of IPv4 for a single TCP port in 3 minutes*
- Leverages 10GbE NICs and PF_RING sockets
- Development continues at GitHub

$ masscan 0.0.0.0/0 -p 80

Nmap

Nmap 6.40 makes scanning mo-better!

- Performance improvements all around
- Tons of new scripts and fingerprints
- XML + NSE output improvements
- Swiss army knife of scanning

Nmap

Nmap is competitive with the right options

- Combine –sS with –PS for one-pass SYN scans
- Set --min-rate and --min-rtt-timeouts
- Limit retries with –-min-retries

Internet Census 2012

Benign botnet used to scan the internet

- Used over 420,000 devices to scan over 730 ports
- Excellent writeup and a whopping 9Tb of data

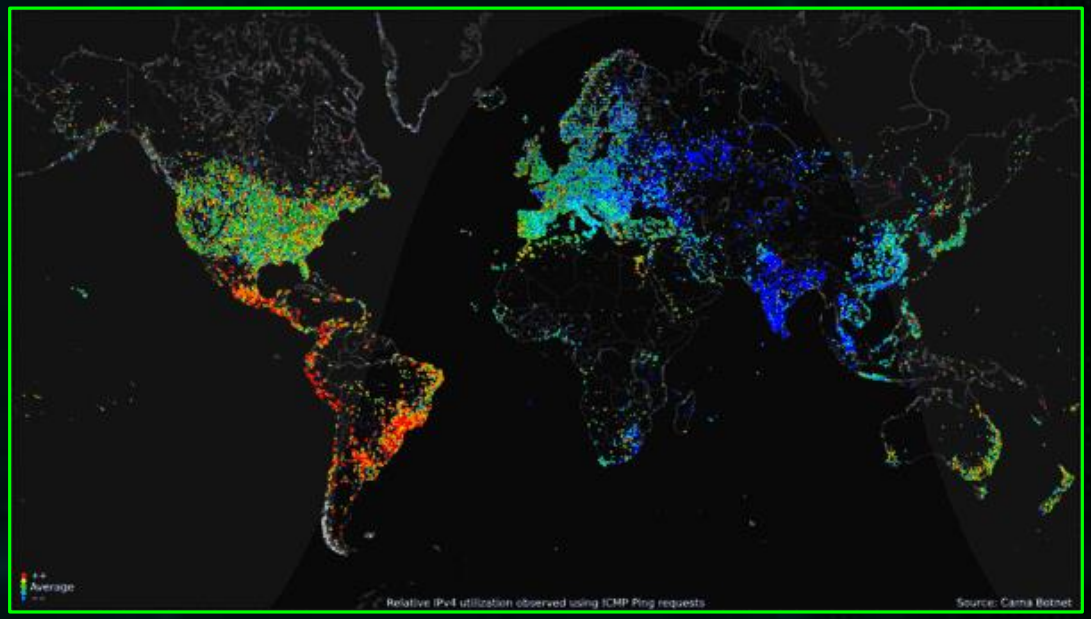

Challenges

Internet scanning has barriers to entry

- Legal concerns vary by region and attitude
- Scans lead to abuse complaints to ISPs
- Computing and time costs

Status Quo

Internet scanning is a niche field

- Challenges prevent widespread adoption
- Value is centered around research
- Businesses can see it as a threat

Internet Scan Data

Internet scan data is incredibly useful

- Identify and quantify widespread vulnerabilities
- Provide due diligence for vendors & partners
- Market share information for products
- Locate unmanaged corporate assets
- Get a handle on shadow IT

Security is Getting Worse

Hard to find any measurable improvement

- Exposures are getting worse each time we look
- VxWorks WDBRPC exposure is increasing
- UPnP has shown minimal improvements
- DDNS DDoS is bad enough
- SNMP is worse

Time for a Change

This is a rock the community can move

- Demonstrate value to IT, security, and the business
- Drive research based on quantified exposure
- Build awareness around public networks
- Hold vendors and ISPs accountable
- Provide ammo for legal reform

zo*hB3LoleLAuNodTwNz7$/:5iUDtsDt1spBu9DtQPEb2comBpuo$NTxBT*kQ9*! Tr0ub4dor&r zo*hB3LoleLAuNodTwNz7$/:5iUDtsDt1spBu9DtQPEb2comBpuo$NTxBT*kQ9*!

Project Sonar

Community project for internet scans

- Open source tools to simplify scanning
- Open datasets for everyone
- Practical applications

http://miniurl.org/sonar

RAPID7

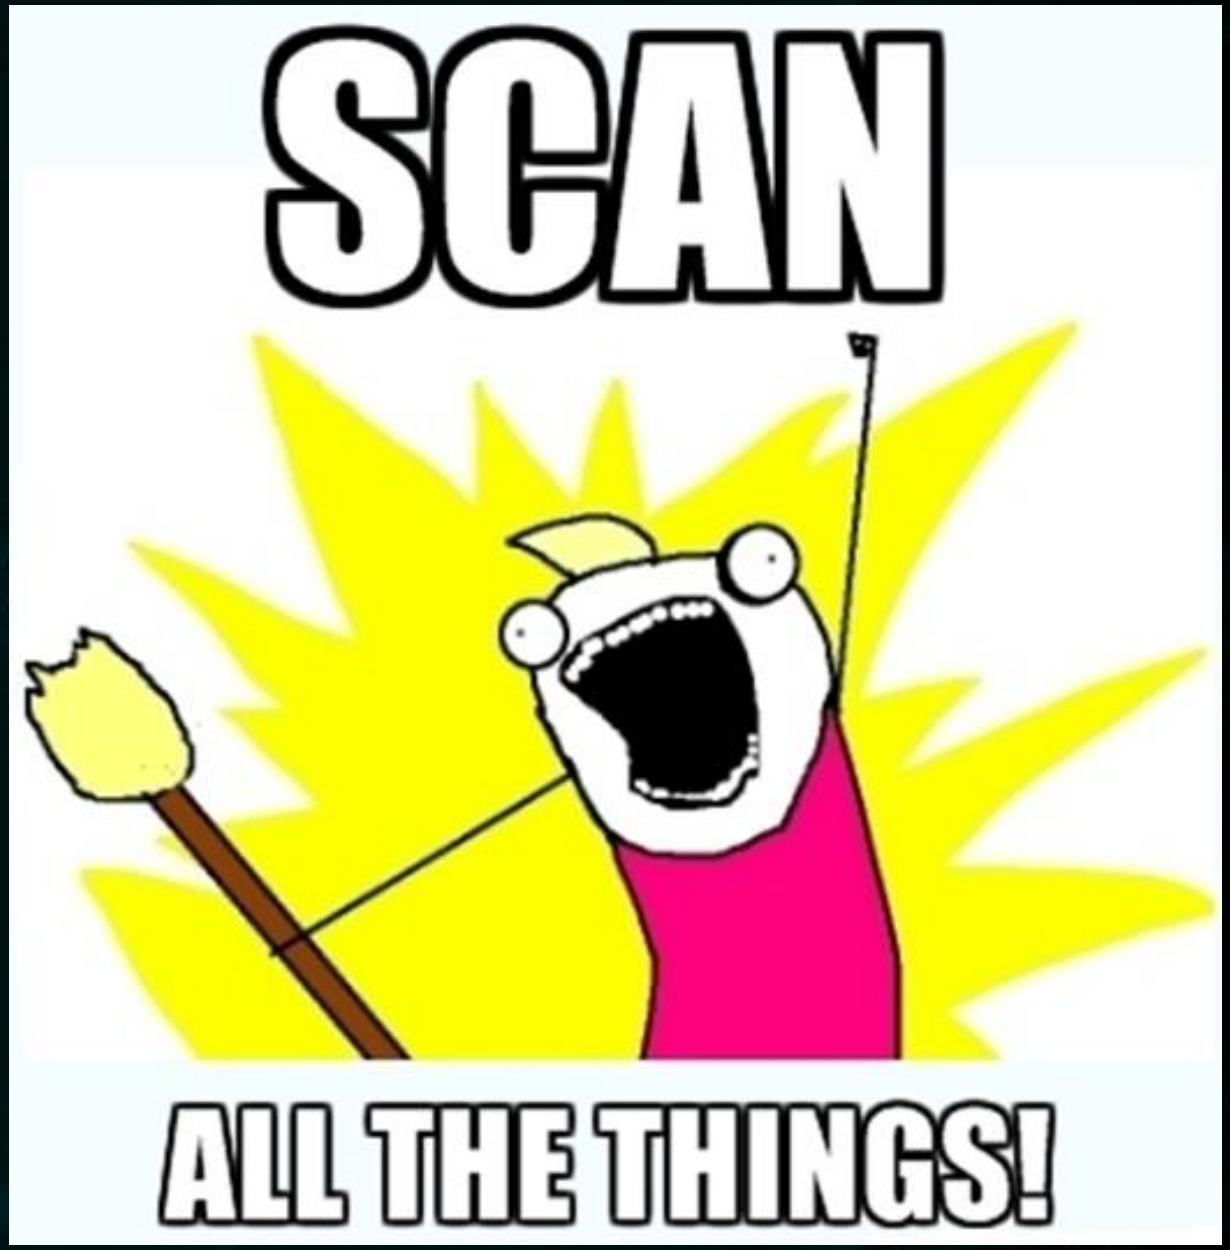

Sonar: Scanning

Integration with existing tools

- UDP probes and processing tools for Zmap
- NSE scripts for running with Nmap
- SSL certificate grabbers
- Fast DNS lookup tools

#ScanAllTheThings

RAPID7

Sonar: Dataset 1

Critical.IO Archive

- Parsed banners across 18 services over 10 months
- Current dataset is in compressed JSON
- Historical view of your networks
- Segmented for easy lookups

#ScanAllTheThings

RAPID7

Sonar: Dataset 1

- 2.4 TB of service fingerprints (355 GB bz2 compressed)
- 1.57 billion records

Management	Email	Discovery	Web
21/tcp	25/tcp	137/udp	80/tcp
22/tcp	110/tcp	1900/udp	443/tcp
23/tcp	143/tcp	5353/udp	8080/tcp
5900/tcp	993/tcp	17185/udp	
3306/tcp	995/tcp		
161/udp			

#ScanAllTheThings

RAPID7

Sonar: Dataset 2

SSL Certificates

- All SSL certs on IPv4 port 443 as of September 10th
- Available as raw certs and parsed IP -> Name pairs
- ~33 million records @ 50 GB (16 GB compressed)
- ~8.6 million unique IP->Name pairs (270 MB)

#ScanAllTheThings

RAPID7

Sonar: Dataset 3

Reverse DNS

- Full reverse DNS for IPv4, regularly updated
- ~1.13 billion records @ 50 GB (3 GB compressed)
- Similar use cases to DeepMagic's PTR search

#ScanAllTheThings

RAPID7

Data Portals & Downloads

ZMap & Rapid7 teams are collaborating

- Launching a shared internet scan data portal
- Accepting data from third-parties (you!)
- Includes all datasets already mentioned
- Also 18 months of SSL scans!

http://scans.io

#GrepAllTheThings?

RAPID7

Examples: Research

You can find zero-day with public datasets

- Easy to identify common vulnerabilities
- Look for min/max and anomalies
- Unix pipelines are all you need

#ScanAllTheThings

RAPID7

Duplicate SSL Certificates

Random things that aren't random

- Any duplicate SSL key is probably a vulnerability
- Tens of thousands of systems with duplicates
- We need eyes to actually classify these
- Identify vendors and report

#ScanAllTheThings

RAPID7

SSL Fingerprinting

SSL certificates make good fingerprints

- Identify all occurrences of an embedded device
- Locate otherwise hard to identify systems
- Enterprise appliances galore

#ScanAllTheThings

RAPID7

Examples: Infosec

Improving your company's security

- Identify external assets you may have missed
- Quickly scan massive networks easily
- Historical data helps with response
- Practical data mining

#ScanAllTheThings

Assets vs Incidents

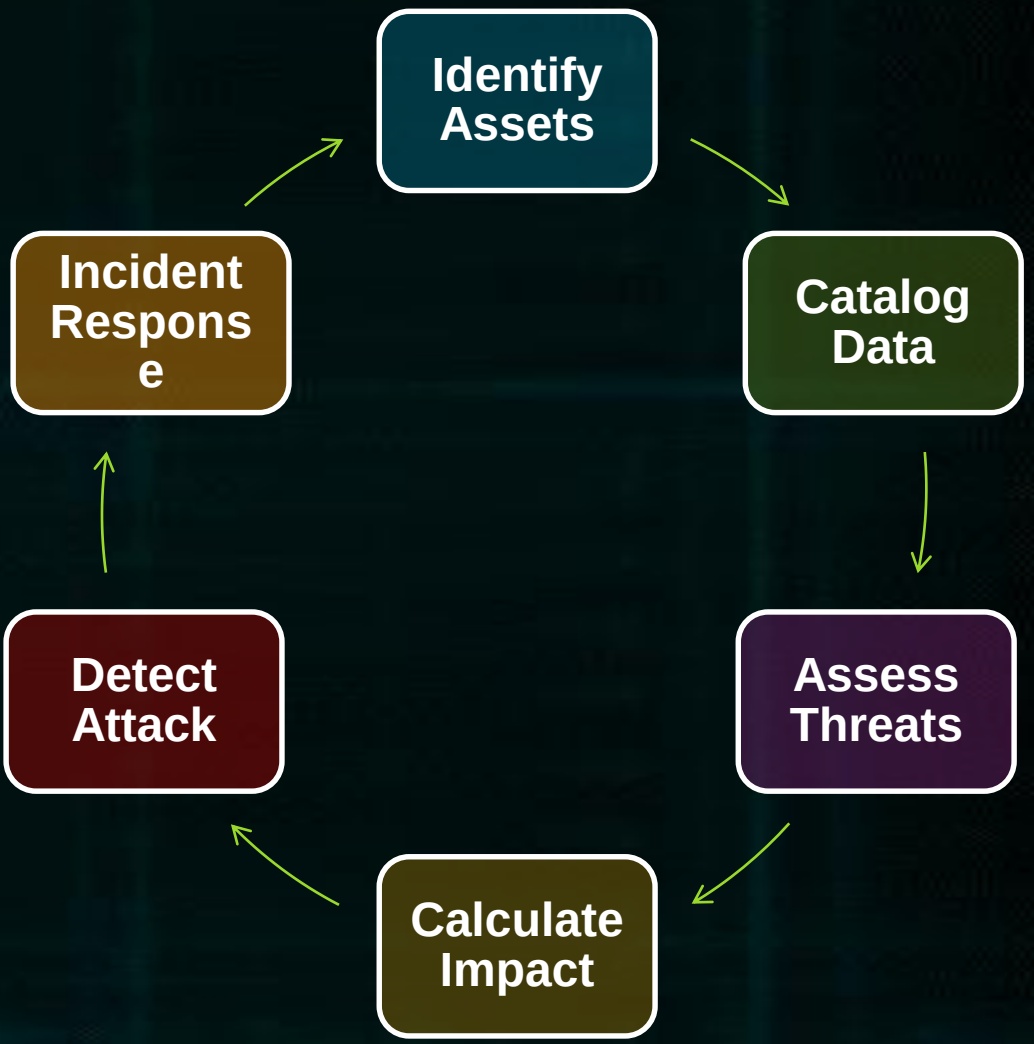

Asset Discovery (SSL)

SSL certificates are ubiquitous

- Every important site has a SSL certificate
- SSL certificates map to domains

Cloud services often use customer certificates

- Identify undocumented third-party services
- May find 10%+ more than your IT knows about

#ScanAllTheThings

RAPID7

Asset Discovery (DNS)

Reverse DNS provides an interesting view

- Forward DNS may not match, but reverse is still set
- Find routers, modems, old ISP connections
- Find VPS services, rogue partners, and VARs
- Accidentally the whole intel agency

#ScanAllTheThings

RAPID7

Quick Risk Assessment

Classify 100,000 nodes in 5 minutes

- Quickly scan a small subset of ports
- Send UDP probes for dangerous services
- Analyze, sort, and prioritize assessment

#ScanAllTheThings

RAPID7

Q & A

http://miniurl.org/sonar

#ScanAllTheThings

RAPID7