DNS as Critical Infrastructure

BARRY IRWIN

BRUCON 0X0E '22

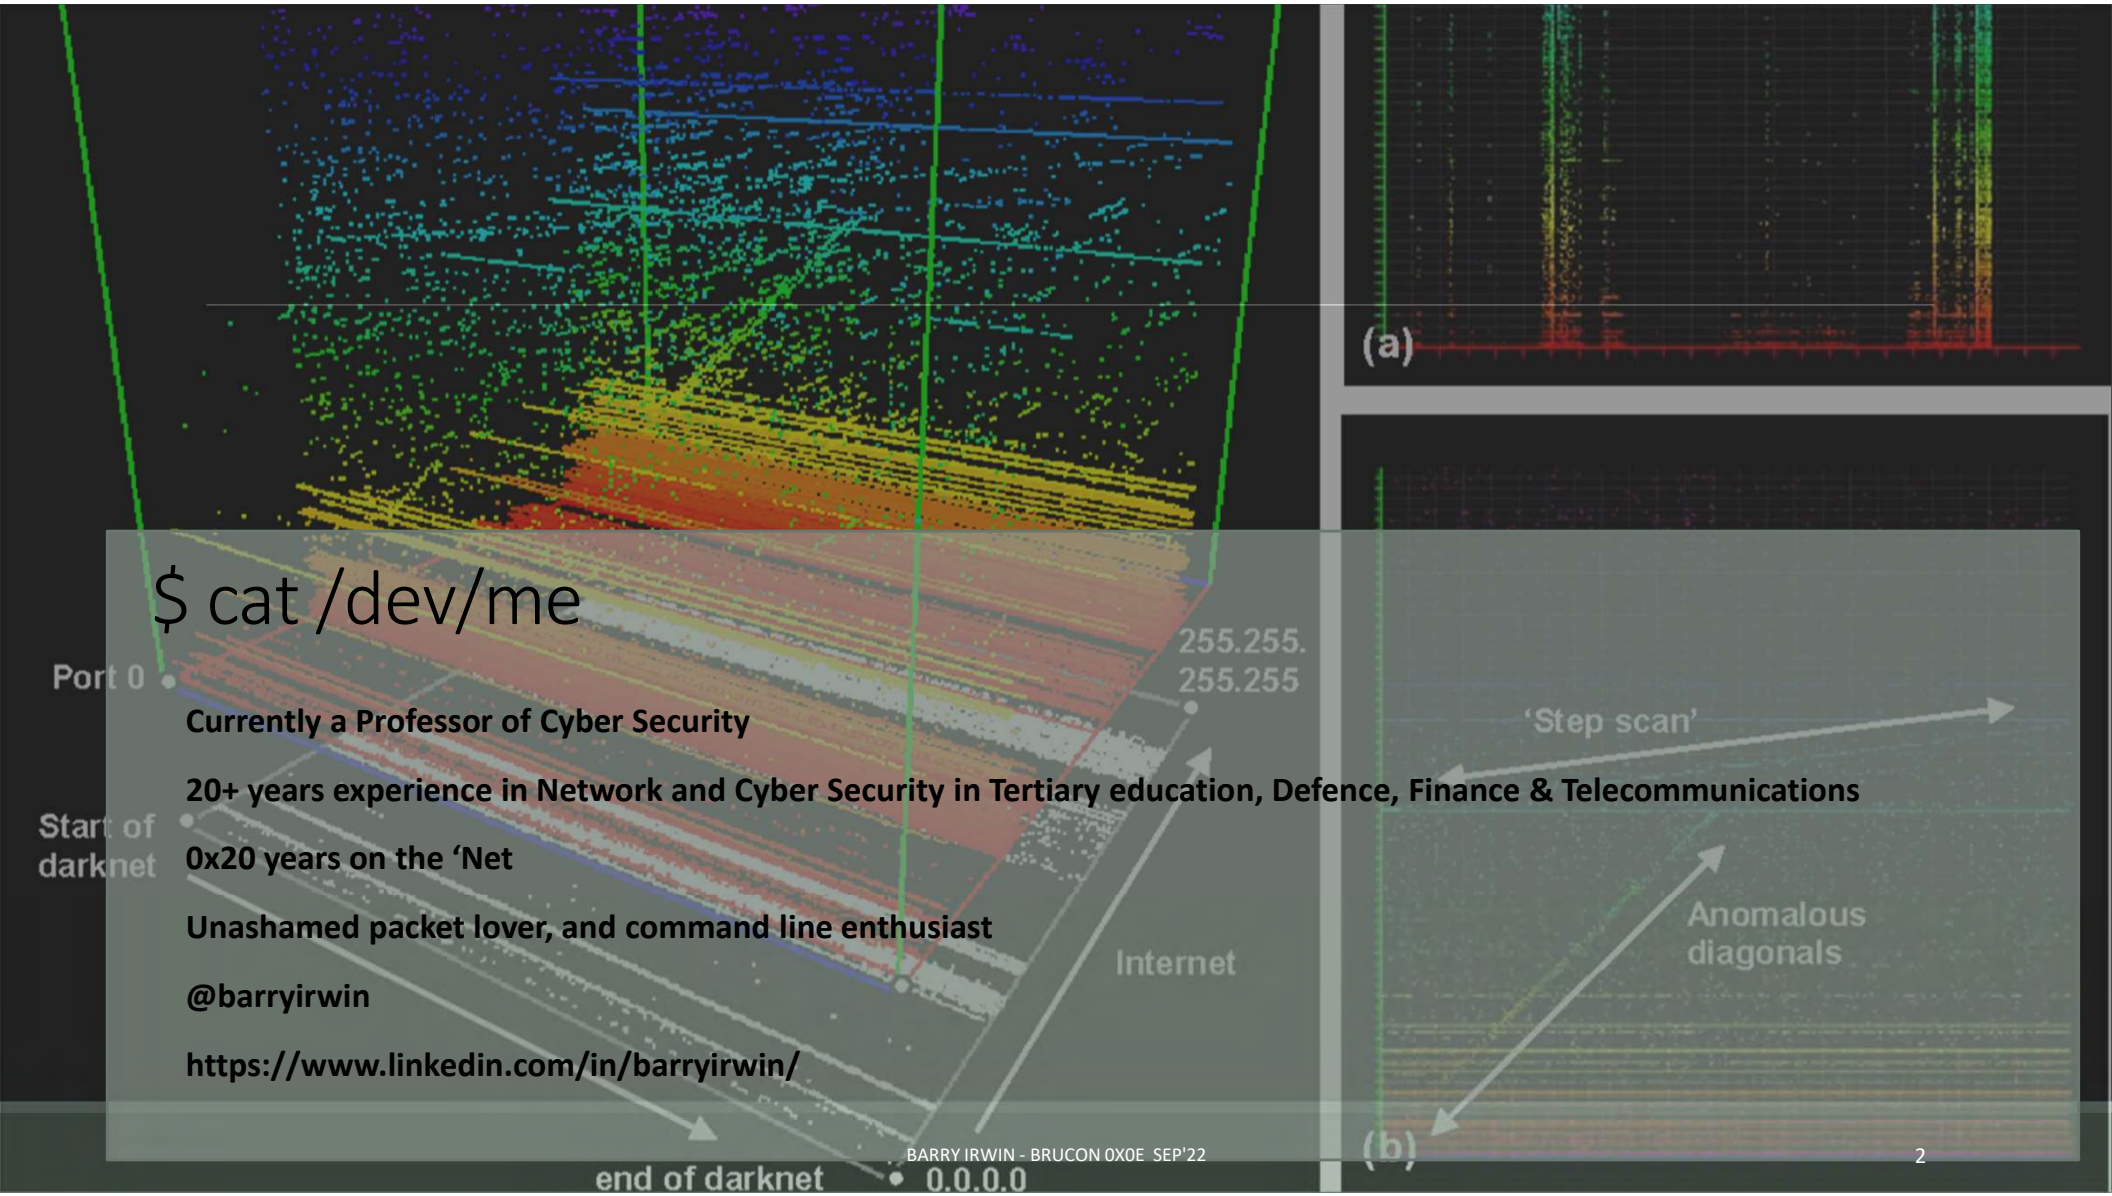

$ cat /dev/me

Currently a Professor of Cyber Security

20+ years experience in Network and Cyber Security in Tertiary education, Defence, Finance & Telecommunications

0x20 years on the 'Net

Unashamed packet lover, and command line enthusiast

@barryirwin

https://www.linkedin.com/in/barryirwin/

Rhykenology

THE STUDY AND COLLECTING OF WOODWORKING PLANES

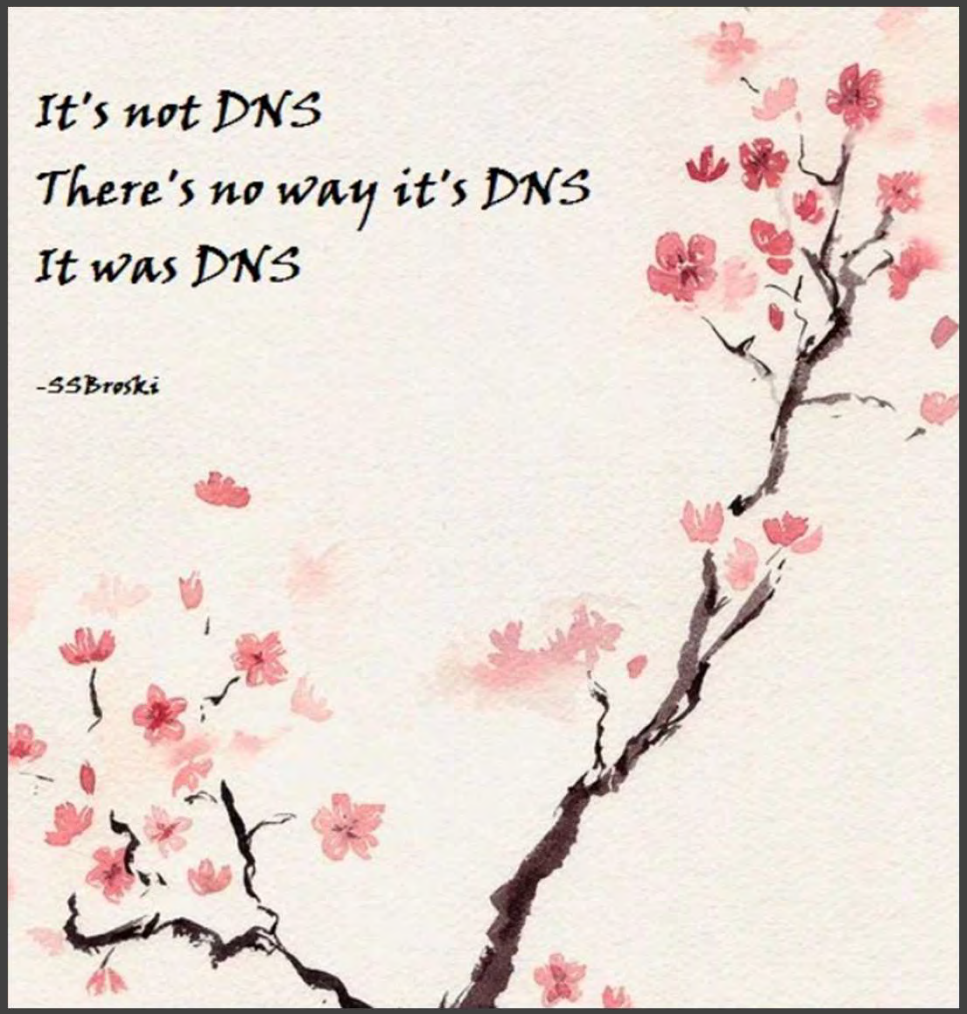

Why DNS?

TL;DR –

Things break,
Badly,
Without DNS

DNS Resiliency

"Ensure DNS Redundancy and High Availability"

Best practice,

- Diversification
- Logical and Geographic distance
- Bind Operators Guide (The BOG)

- Pretty much everything we have today relies on DNS **and** on DNS being functional in terms of providing resilience/loadbalanching/functional service

Disclaimer

The results here are ‘broad strokes’

Details are blinded to protect the (potentially) vulnerable

This is based on a series of snapshots over a period of months

Results are largely constrained by the accuracy and representation of the input data (getting good inputs is a challenge)

No hard, concrete solutions, just some concerned flag waving (and ideas)!

Interpretation and views are **my own**

The experiment?

What is the diversity of the ccTLDs??

What proportion is hosted in *vs*. outside $cctld?

What is the risk to DNA as critical infrastructure?

What is the adoption of DNSSEC?

What is the degree of adoption of Newer DNSRR's like CAA ?

https://commons.wikimedia.org/wiki/File:Louis_Lobera_d%27Avila_in_his_study_woodcut_by_H._Burgkmair.jpg


© Kacper Pempel/REUTERS

Gathering Data

How do we begin?

- Need domain lists (AXFER doesn't work ☺)
- Harder than one would think
  - Built up from various online sources and lists
  - Domain authorities don't want to share data …
    … because bad things (tm)
- These are imperfect:
  - Hostnames != Domains
  - ccTLDS have different approaches.
  - Approx. 8% on average NXDOMAIN
  - Timeouts/Refused <1% after 3 runs
  - Runs over last 5 months have shown to be fairly consistent

Domains of interest

ccTLD with no finite 2nd level structure - .no .be .ru

ccTLD - ‘commercial 2nd level’

- .uk – co.uk used as largest viable proxy
- .au – com.au used
- .za – co.za used

Majestic Million ‘global benchmark’ (??)

Issues with processing:

- Timeouts
- RFC1918 DNS servers
- NXDomains
- Refused <0.001%

The final ~~treasure~~^W data

com.au	1627814
co.uk	4405918
co.za	948326
ru*	243258
no	570869
be	1203248
Majestic	1000000

CAVEAT:

- Data is volatile and domains expire and new ones registered.
- Imperfect is better than nothing
- Snapshots are not 100% accurate
- ~10 million domains, 320K NS

There are three kinds of lies — lies, damned lies and statistics.

(Mark Twain)

Lies, damned lies, and statistics.

(Benjamin Disraeli)

Data quality ?

IS INCOMPLETE DATA, BAD DATA ?

HOW DOES IT COMPARE TO **NO** DATA ?

Processing...

Cleaning, and more cleaning

Hack it fast… is slow!
- For *x* in domainlist do; ~3 days/600K
- Python – 18h / Million

Optimised approach using *zdns*
- ~~~38 minutes /600K domains~~~
- ~42 minutes / Million domains (NS)
- ~12minutes / Million (DS)
- ~30 minutes/ Million (CAA)

Could be further optimised given consideration of structure & distribution of domains
- Caching
- Parallel Processing
- Need to manage limits and 'be nice' to servers
  - 1500 QPS seems about as high as is reasonable, but 'it depends'

Processing Challenges

Raw inputs collected were not ‘clean’ (surprise!)

Issues to consider when running collection

- Expired domains
- Upstream routing
- Timeouts
- Configuration errors (surprising number of RFC1819/3330 addresses exposed)
- Try, and try again

Run data needed post processing

- Record what worked
- Prune NX
- Retry Timeouts/refused/SERVFAIL

Tools

Bash

sed / awk et al.

jq (use modern data without the pain)

ZDNS

Netcat

Team Cymru and Maxmind for Geolocation

Some spreadsheets (for checking ;^>)

Approach

Iterative

Collect as much as possible (within reason)

Batched collection (help with caching)

Scaling

~72 hour window for retries

Work out what data **is** important

Deal with massive data explosion – JSON and flat files….

JSON provides unintended benefits ito additional data (efficiency)

Maybe a RDBMS would be the better approach

Act I – Adoption rates

WHAT IS THE ADOPTION OF DNSSEC AND CAA LIKE?

TL;DR – ☹

Doveryai, no proveryai
(Доверяй, но проверяй)
- Russian proverb

Trust, but verify
- Ronald Regan

DNSSEC

PROVIDES CRYPTOGRAPHIC AUTHENTICATION OF DATA, AUTHENTICATED DENIAL OF EXISTENCE, AND DATA INTEGRITY, BUT NOT AVAILABILITY OR CONFIDENTIALITY

It’s a sad state of affairs.

1999 – RFC 2065/2535 is the birth of DNSSEC

2005 – RFC 4033/4/5 – DNS is ready for Prime time, RIPE starts deployment

2010 - .org is first TLD to be signed. Followed by root zone.

2013 - More than 100 ccTLDs and all legacy TLDs signed, **all** new TLDs required to be signed.

Now nearly another 10 years on… **Generally poor adoption observed**

Some counties are higher than others NO, SE, NL are >50% (based on other research)

	Chung et al (2017)		Roth et al. (2019)	
TLD	Domains	Signed domains	Domains	Signed domains
.com	118,147,199	0.7%	140,438,505	0.8%
.net	13,773,903	1.0%	13,408,301	1.1%
.org	9,682,750	1.1%	10,066,388	1.1%
.NL	5,674,208	51.6%	5,860,418	54.1%
.SE	1,388,372	46.7%	1,450,441	56.9%

DNSSEC adoption rates

Domain	Tested	Have	%
Majestic	996338	3313	0,33
com.au	1627814	5735	0,35
co.uk	2656362	59238	2,23
be	1064328	291691	27,41
co.za	948326	1509	0,16

You can't trust code that you did not totally create yourself.

- Ken Thompson

CAA

DNS RECORD USED TO PROVIDE ADDITIONAL CONFIRMATION FOR THE CERTIFICATION AUTHORITY (CA) WHEN VALIDATING AN SSL CERTIFICATE

Certification Authority Authorization

2010 – First published

2019 – RFC8659 is the latest standard

Intended to provide explicit statement of CA's permission to issue certificates for a domain

- **Issue** - authorizes the CA specified to issue certificates for the domain
- **Issuewild** – like issue but takes priority regarding wildcard certificates.
- **Iodef** – specific contact method to report invalid certificate requests

A relatively new protocol, but still low adoption rates.

How does this provide security ?

Top 10 from the Majestic

13217	issue	letsencrypt.org	
9212	issue	comodoca.com	
8890	issuewild	letsencrypt.org	
8052	issuewild	comodoca.com	
7357	issue	digicert.com;	cansignhttpexchanges=yes
7255	issuewild	digicert.com;	cansignhttpexchanges=yes
7012	issue	pki.goog;	cansignhttpexchanges=yes
6993	issuewild	pki.goog;	cansignhttpexchanges=yes
5520	issue	amazon.com	
5294	issue	digicert.com	

‘Trusted’ CA’s – Majestic (top 12)

Rank	%	CA
1	19,19	digicert.com
2	19,16	letsencrypt.org
3	14,96	comodoca.com
4	12,99	pki.goog
5	6,67	amazon.com
6	4,12	globalsign.com
7	4,10	sectigo.com
8	3,63	amazonaws.com
9	3,16	amazontrust.com
10	3,07	awstrust.com
11	2,02	godaddy.com
12	0,73	entrust.net
	93,80	
		N=115436

CAA Adoption rates

Domain	Tested	Have	%
Majestic	916904	34208	3,73
com.au	1627814	13781	0,85
co.uk	2656362	20566	0,77
be	1064328	11165	1,05
co.za	948326	3775	0,40

Act II – Critical Risks

HERE BE DRAGONS.
ITS 2 AM .

DO YOU KNOW WHO CONTROLS YOUR DNS?

WHAT NATION-STATE IS GOING TO RUIN YOUR DAY ?

Finding value

Huge amounts of data (36GB) to deal with….

- ..this only scrapes the surface of what can be found

"Premature Optimization is the root of all evil" (Knuth)

- Of the domains surveyed, **all** are at risk of influence by foreign players impacting DNS
- There is a wealth of opportunity for further exploration
- Threat modelling for DNS ?

Australia (com.au)

N=26217

Australia controls 20% of the Name servers used

CN (91) and RU (61) Servers

NZ hosts 148

Issues with geographic isolation

67% North America

7% Western Europe

Rank	CC	#NS	%of total
1	US	16389	63%
2	AU	5244	20%
3	CA	944	4%
4	DE	563	2%
5	FR	491	2%
6	GB	341	1%
7	NL	320	1%
8	IN	182	1%
9	SE	148	1%
10	NZ	148	1%
	Total		*94%*

South Africa (co.za)

N=18197

South Africa controls 14% of the Name servers used

CN (55) and RU (56) Servers

BW, MZ, ZW, LS <10 servers

Issues with geographic isolation

63% North America

19% Western Europe

Rank	CC	#NS	%of total
1	US	11326	62%
2	ZA	2524	14%
3	DE	1283	7%
4	FR	667	4%
5	GB	433	2%
6	NL	384	2%
7	CA	219	1%
8	BG	140	1%
9	AU	121	1%
10	CH	110	1%
	Total		95%

United Kingdom (co.uk)

N=68614

UK controls 20% of the Name servers used

CN (142) RU (298) IR (54) Servers

Issues with geographic isolation

43% North America

23% Western Europe (low risk)

Rank	CC	#NS	%of total
1	US	28238	41%
2	GB	14003	20%
3	DE	6385	9%
4	FR	4194	6%
5	NL	2608	4%
6	CA	1557	2%
7	SE	1312	2%
8	BG	781	1%
9	IT	700	1%
10	TR	681	1%
	Total		88%

Belgium (.be)

N=32569

Belgium controls **4%** of the Name servers used

CN (104) RU (186) IR (8) BY (5) Servers

Issues with geographic isolation

38% North America

47% Western Europe (low risk)

Rank	CC	#NS	%of total
1	US	11965	37%
2	NL	6386	20%
3	DE	3596	11%
4	FR	3482	11%
5	BE	1411	4%
6	GB	725	2%
7	CA	539	2%
8	CH	479	1%
9	SE	394	1%
10	IT	374	1%
	Total		90%

Norway (.no)

N=16027

Norway controls 6% of the Name servers used

CN (90) RU (82) IR (2) Servers

SE and DK have 12%

Issues with geographic isolation

56% North America

20% Scandinavia

29% Western Europe

Rank	CC	#NS	%of total
1	US	8794	55%
2	SE	1840	11%
3	NO	1037	6%
4	DE	924	6%
5	FR	788	5%
6	NL	453	3%
7	GB	308	2%
8	FI	149	1%
9	CA	149	1%
10	DK	136	1%
	Total		91%

Russian Federation (.ru)

N=17050

Russia controls 45% of the Name servers used

CN (49) BY (40) Servers

UA 223 Servers

Issues with geographic isolation

49% Western Europe and USA

Rank	CC	#NS	%of total
1	RU	7651	45%
2	US	5661	33%
3	DE	1114	7%
4	FR	393	2%
5	NL	353	2%
6	GB	225	1%
7	UA	223	1%
8	CZ	203	1%
9	EE	92	1%
10	BG	91	1%
	Total		94%

Majestic Million

N=140444

200 countries

Rank	CC	#NS	%of total
1	US	56062	40%
2	DE	10054	7%
3	FR	6446	5%
4	RU	5352	4%
5	JP	5307	4%
6	GB	4622	3%
7	CA	4375	3%
8	CN	4197	3%
9	NL	4071	3%
10	ES	2181	2%
	Total		73%

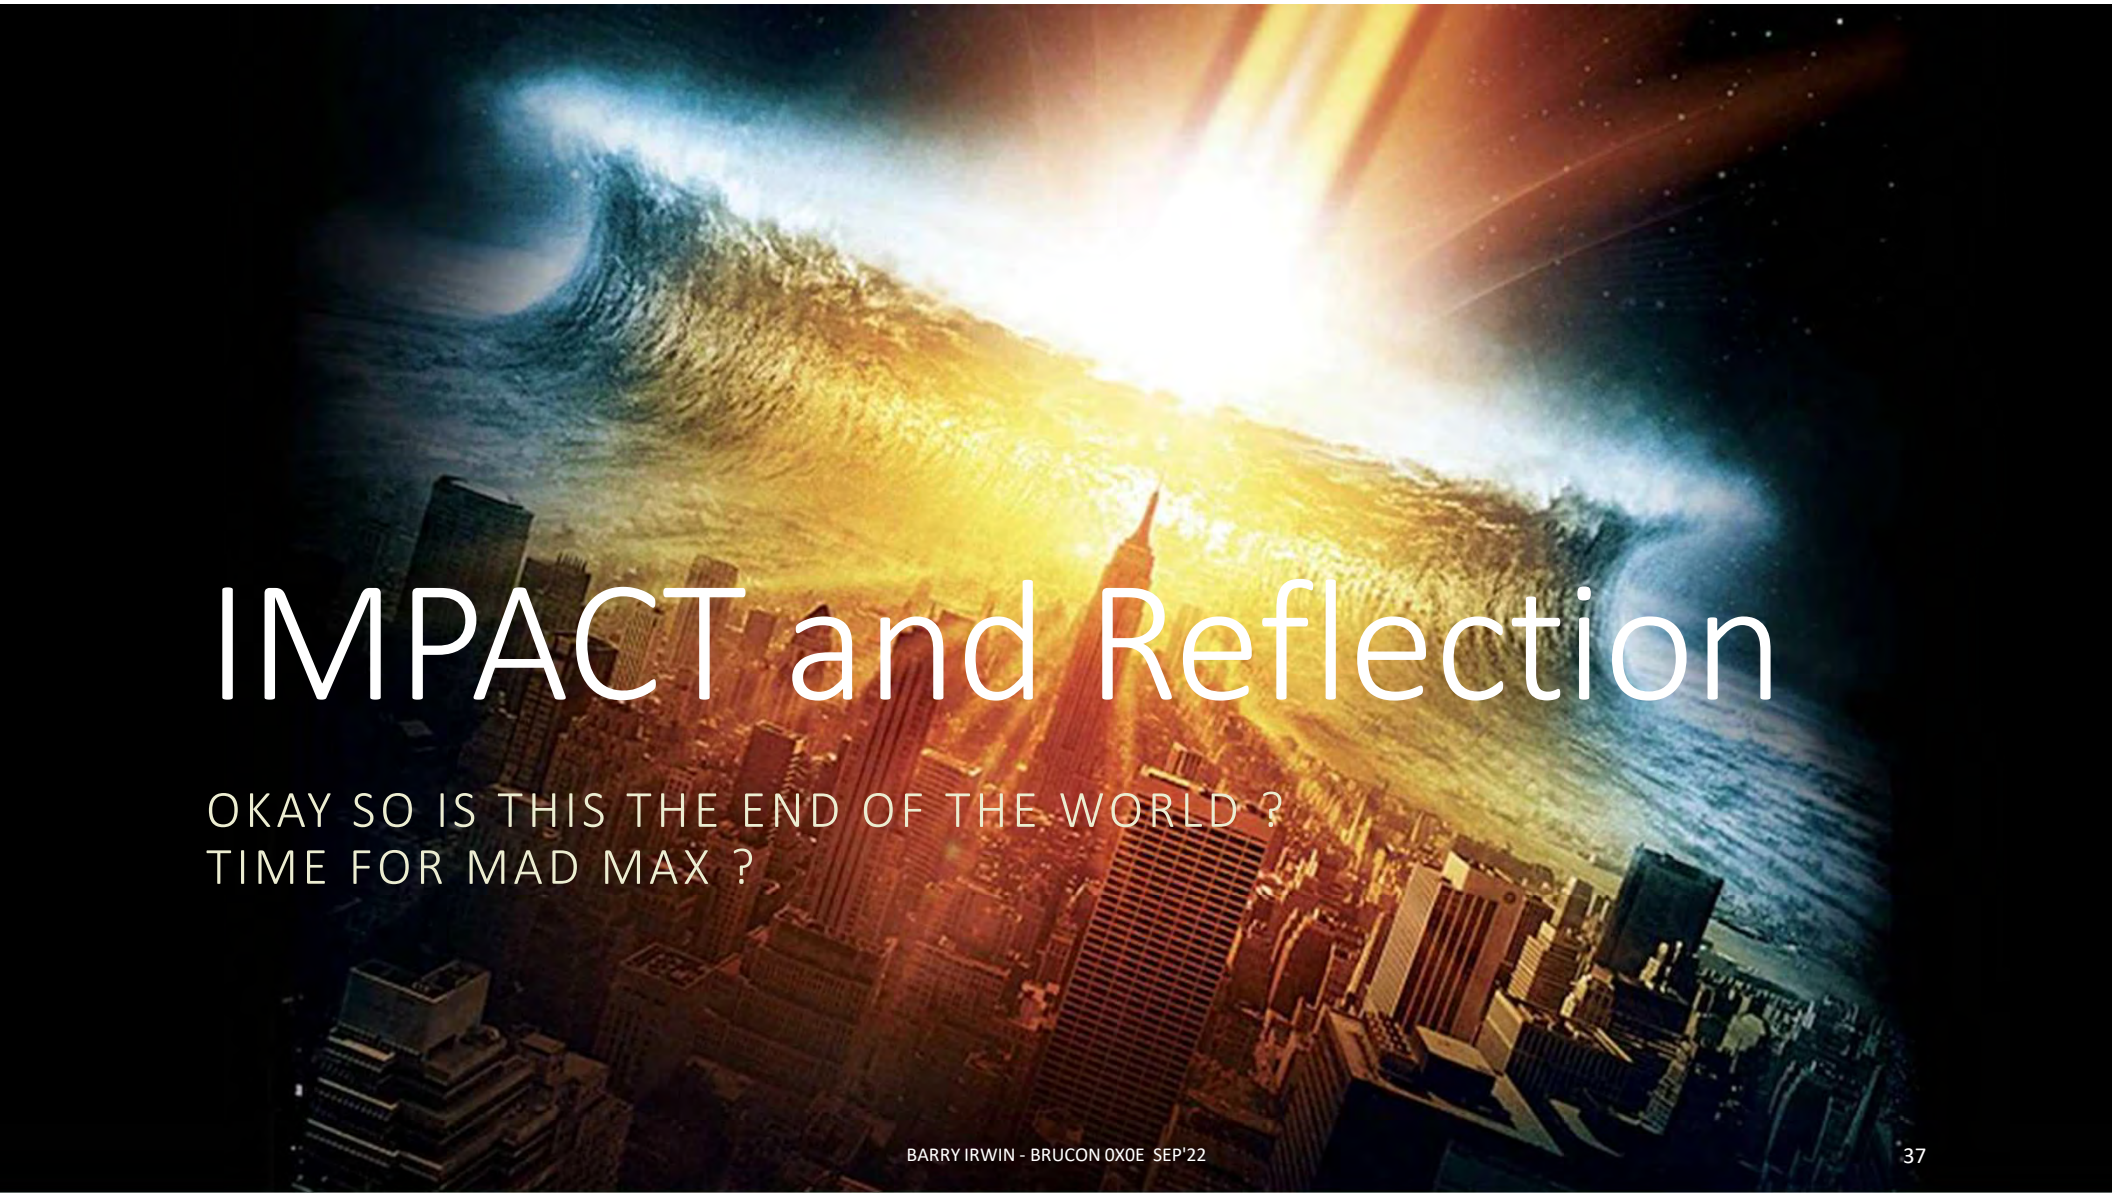

IMPACT and Reflection

OKAY SO IS THIS THE END OF THE WORLD ?
TIME FOR MAD MAX ?

"The supreme art of war is to subdue the enemy without fighting."
— **Sun Tzu, The Art of War**

Impact I

ATTACK ON <20 IP ADDRESSES COULD RENDER ~75 % OF NORWEGIAN CCTLD'S UNWORKABLE.

"Victorious warriors win first and then go to war"
— Sun Tzu, The Art of War

Impact II

ATTACK ON TOP 5 UK NS PROVIDERS RENDERS 10% OF CO.UK AND ~440K DOMAINS UNWORKABLE.

The data shows that this is most likely a hundreds-of-thousands to millions of victims issue.
- Dan Kaminsky on DNS flaws

Impact III

ATTACK ON TOP 5 NS PROVIDERS FOR .BE COULD RENDER 20% OF DOMAINS UNWORKABLE.

We seem to be our own worst enemies. We should require critical U.S. infrastructure to remain in U.S. hands.
— DL Hunter, US politician

Impact IV

BIG DNS PROVIDERS HAVE RESILIENCE.

MOST SMALLER ONES DO NOT.

All IP addresses are equal,
but some are more equal

- N4pol30n && 5now|3a11

Impact V

THERE ARE SOME PORTIONS OF IPV4 ADDRESS SPACE THAT SHOULD BE CONSIDERED MORE IMPORTANT THAN OTHERS.

SERVERS CAN BE RELOCATED – ONLY BECAUSE OF DNS. DNS IS HARD(ER)!

Impact?

Relatively small number of systems being targeted could result in out of scale impact

Risk of foreign hosted systems ?

Is this significant, or do stats mislead ?

DNS servers as critical Infrastructure ?

What is the impact of having foreign hosted domains ?

Do we know what we don't know ?

Threat modelling guides for DNS ?

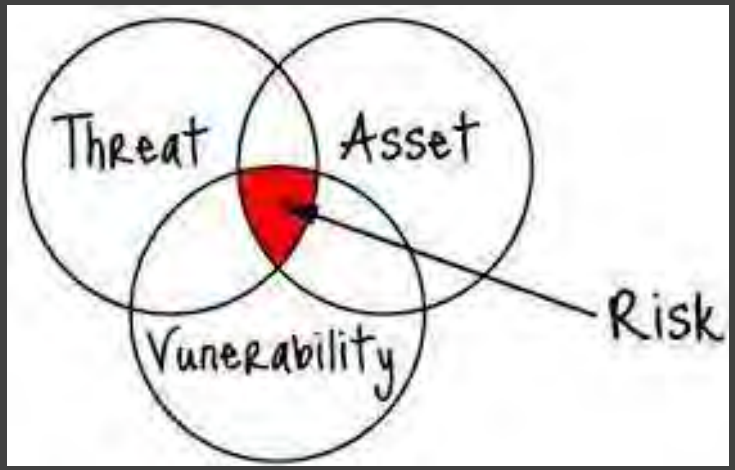

What happens when the unexpected occurs ?

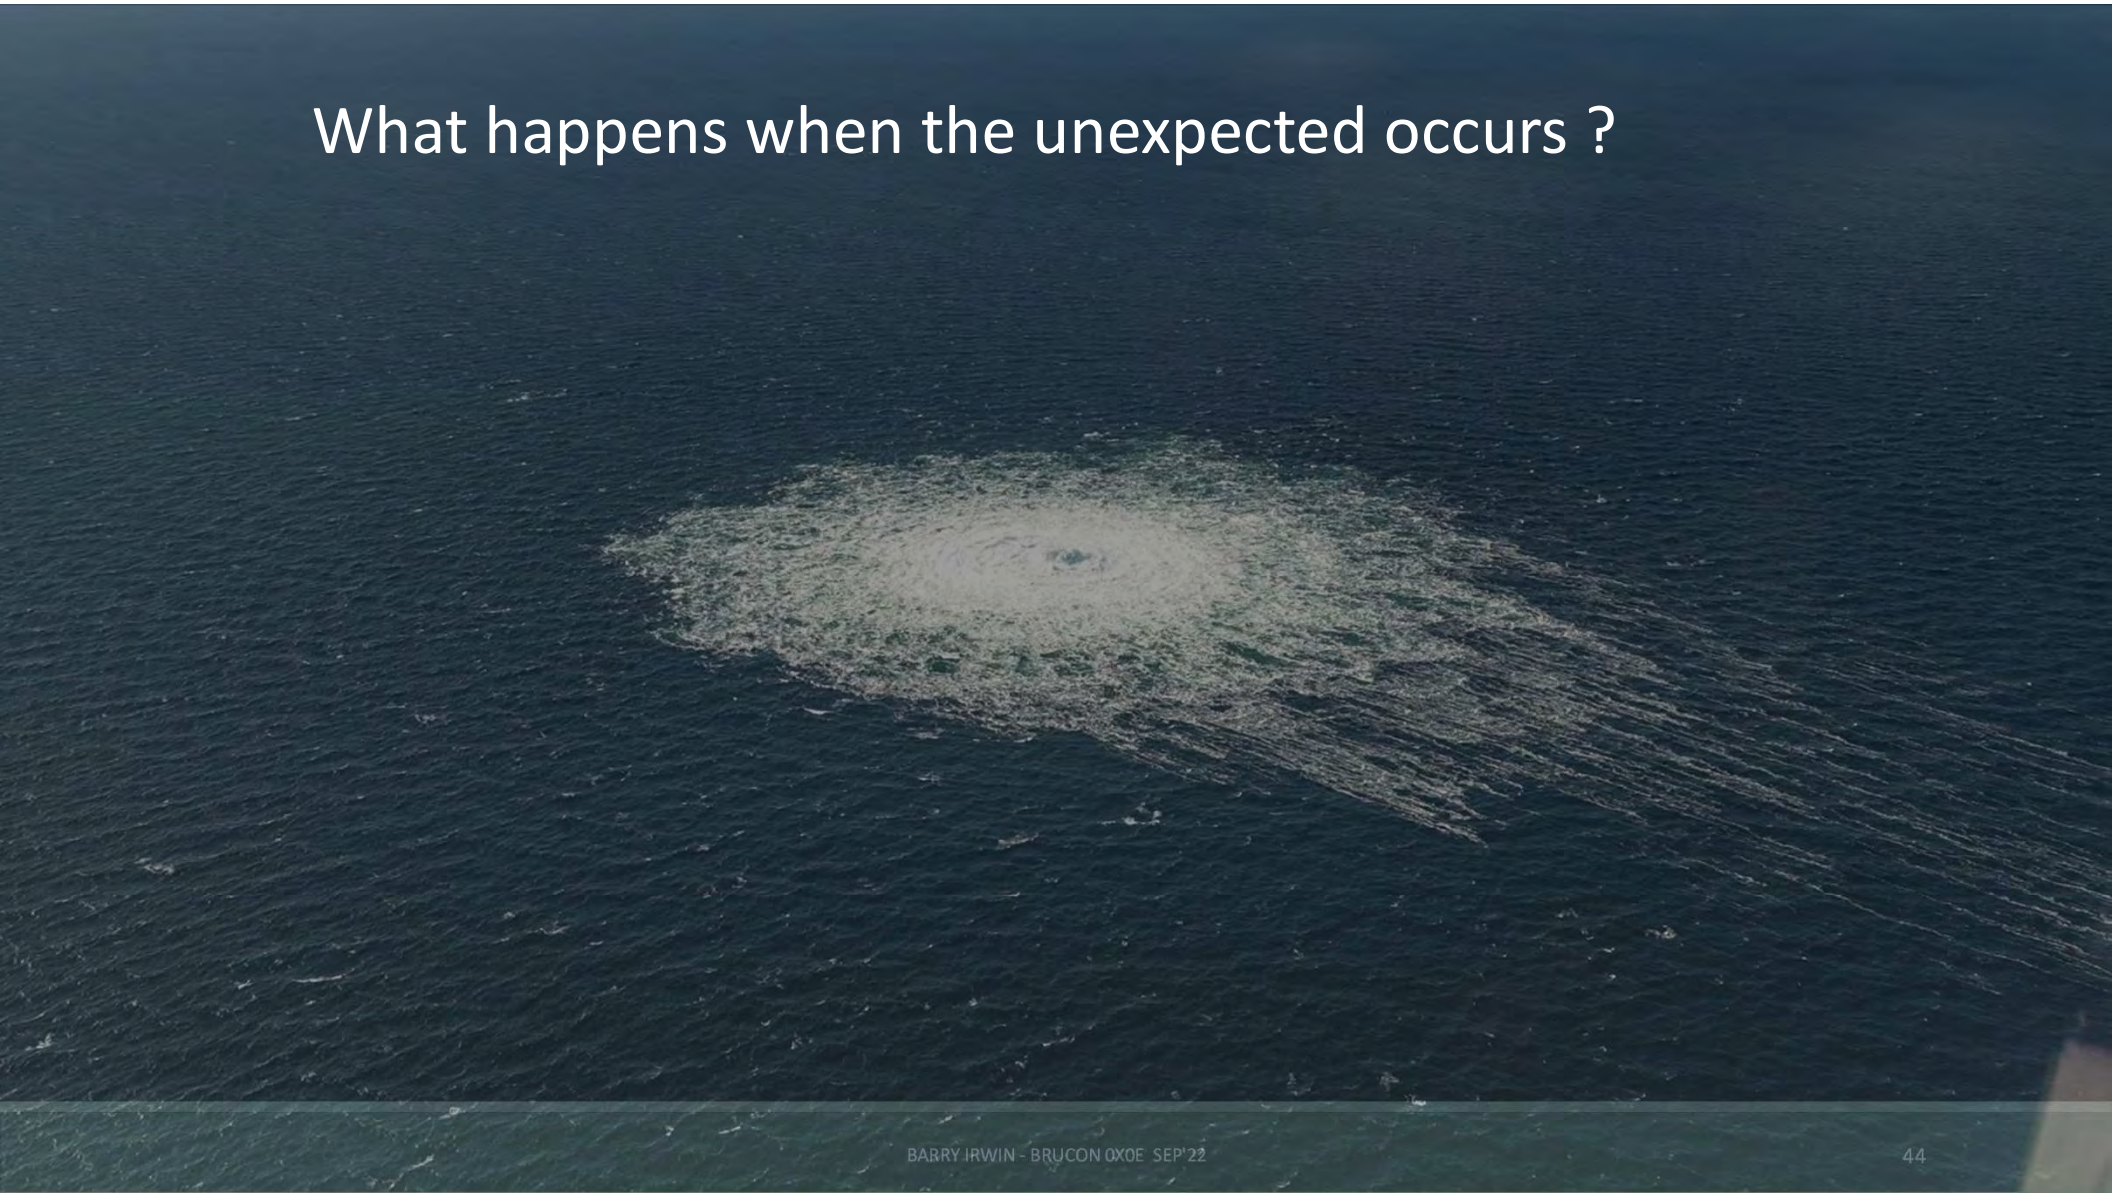

The Devil is in the Details

I WAS INTERESTED IN IMPLEMENTS OF MASS DESTRUCTION (FROM AN ACADEMIC POINT OF VIEW).

DAN FARMER

Complex problems..

DNS is an amazing technology

Surprisingly poorly understood

No-one cares when it works

Arguably the world largest dynamic distributed datastore

Distributed Nature makes it hard to create momentum for change ?

- Care, Coordination, Competency

Are all domains (and sub domains) equally important ?

DNS as a backbone for trust?

Things learned &Things to do

Is there a Problem ?

How bad is it ?

Is it Really bad ?

Should one worry ?

How to make it better?

- Work with CSIRTS, National registrars
- Awareness

Longer term monitoring needed

There are more questions now than when the work started!

Barry Irwin @barryirwin

IF YOU ARE INTERESTED TO KNOW MORE, COME SAY HELLO!

ESPECIALLY (NATIONAL) CSIRTS/ REGISTRARS/ RESEARCHERS