Decentralising Data Collection and Anonymisation

Steven Murdoch

University College London

An anonymisation system has architecture, security definition and mechanism

Architecture

How is data collected, stored, processed and distributed?

Security definition

What security and privacy properties must the system preserve, and what level of confidence is required?

Mechanism

How is access to data restricted or data altered in order to meet security definition?

In this presentation I'm going to show an alternative architecture to the traditional centralized data curator allowing the relaxation of the "plausible deniability" security property

A trusted, centralised architecture is common but creates risk and expense

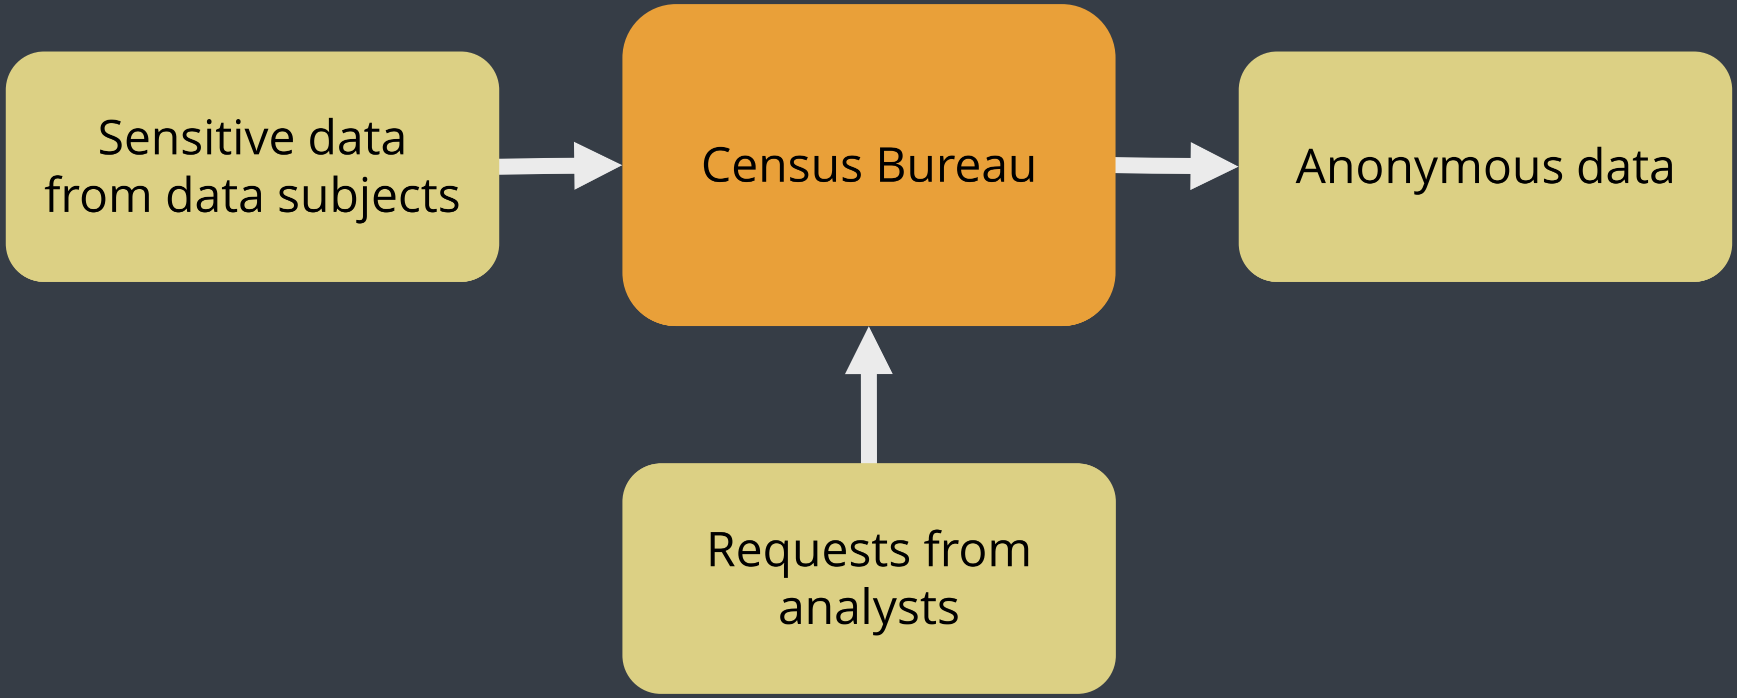

Sensitive data from data subjects

Census Bureau

Anonymous data

Requests from analysts

A trusted, centralised architecture is common but creates risk and expense

Trusted

In a position to violate security properties

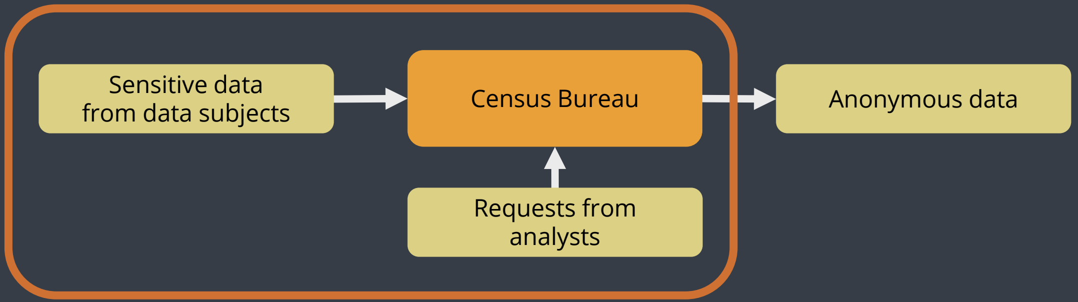

Data submitted to central authority is inherently sensitive

To give data subjects confidence that their data will be protected, strong procedures and technical controls need to be in place, which is expensive in terms of effort, lost opportunity, and monetary cost

A t
con

DEPARTMENT OF DEFENSE STANDARD
DEPARTMENT OF DEFENSE TRUSTED COMPUTER SYSTEM EVALUATION CRITERIA
DECEMBER 1985
DEPARTMENT OF DEFENSE PASSWORD MANAGEMENT GUIDELINE
CSC-STD-002-85
COMPUTER SECURITY REQUIREMENTS
GUIDANCE FOR APPLYING THE DEPARTMENT OF DEFENSE TRUSTED COMPUTER SYSTEM EVALUATION CRITERIA IN SPECIFIC ENVIRONMENTS
CSC-STD-003-85
TECHNICAL RATIONALE BEHIND CSC-STD-003-85: COMPUTER SECURITY REQUIREMENTS
GUIDANCE FOR APPLYING THE DEPARTMENT OF DEFENSE TRUSTED COMPUTER SYSTEM EVALUATION CRITERIA IN SPECIFIC ENVIRONMENTS
CSC-STD-004-85
25 June
NATIONAL COMPUTER SECURITY CENTER
NCSC-TG-001 VERSION-2
A GUIDE TO UNDERSTANDING AUDIT IN TRUSTED SYSTEMS
1 June 1988
NCSC-TG-002 VERSION-1
NCSC-TG-003 VERSION-1
NCSC-TG-004 VERSION-1
NCSC-TG-005 VERSION-1
NCSC-TG-006 VERSION-1
A GUIDE TO UNDERSTANDING CONFIGURATION MANAGEMENT IN TRUSTED SYSTEMS
NCSC-TG-007 VERSION-1
NCSC-TG-008 VERSION-1
NCSC-TG-009 VERSION-1
NCSC-TG-010 VERSION-1
NCSC-TG-011 VERSION-1
NCSC-TG-013 VERSION-1
NCSC-TG-014 VERSION-1
Guidelines for Formal Verification Systems
NCSC-TG-015 VERSION-1
NCSC-TG-016 VERSION-1
NCSC-TG-017 VERSION-1
NCSC-TG-019 VERSION-1
NCSC-TG-020-A VERSION-1
NCSC-TG-021 VERSION-1
NCSC-TG-022 VERSION-1
NCSC-TG-023 VERSION-1
NCSC-TG-024 VOLUME 1/4 VERSION-1
NCSC-TG-024 VOLUME 2/4 VERSION-1
NCSC-TG-025 VERSION-2
NCSC-TG-026 VERSION-1
NCSC-TG-027 VERSION-1
NCSC-TG-029 VERSION-1
NCSC-TG-030 VERSION-1

D

T
pr
ter

ve in

Within the centralised architecture, security definition and mechanism may vary

Data released could be noninteractive (anonymised version of original data, aggregate statistics), or interactive (results of a query interface)

		Security definition	
		Differential privacy	***k*-anonymity**
Architectural variant	**Interactive**	Adding Laplace noise to queries	Generalisation within query
	Non-interactive	Aggregate statistics	Record suppression

> “To me, differential privacy may be as authoritarian in its conceptual underpinnings as [Identity Based Cryptography]”

The Moral Character of Cryptographic Work , Phillip Rogaway

Architecture, security definition and mechanism are linked but often conflated. Authoritarian tendencies come from trusted central authority, not security definition. Criticism above applies mainly to interactive centralized architecture

Confidence in institutions is at an all-time historical low

% who trust the govt in Washington always or most of the time

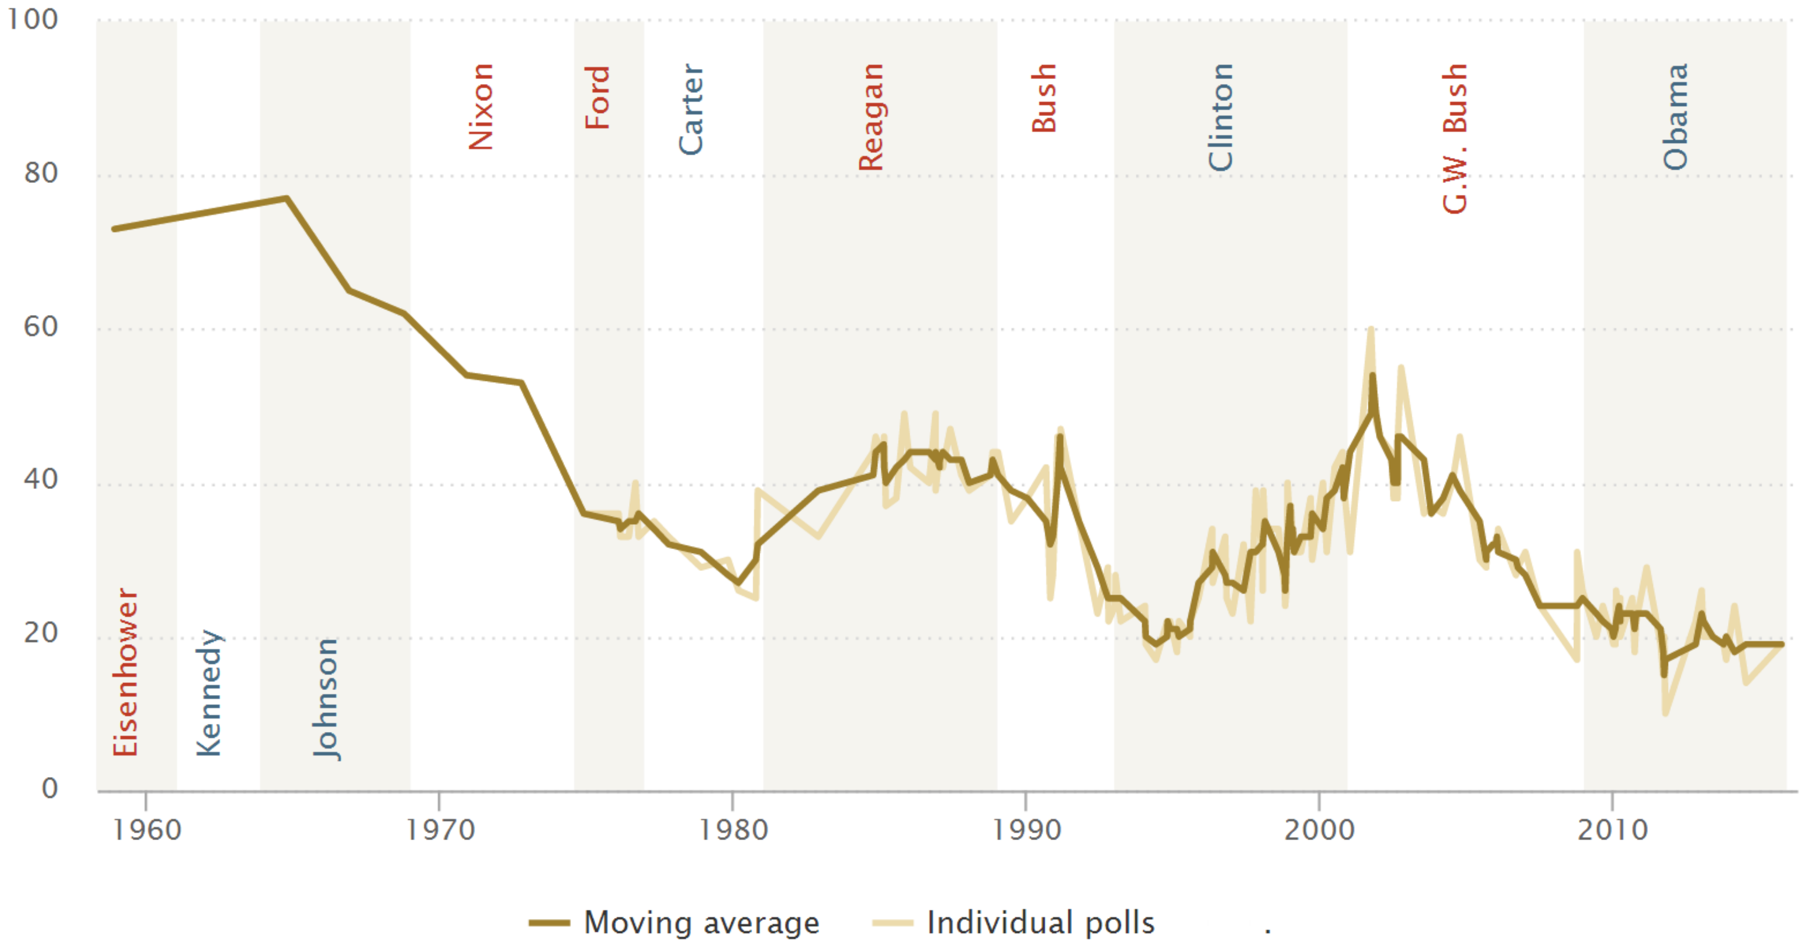

Pew Research Center

If data is anonymised by data subject, the central authority is no longer trusted

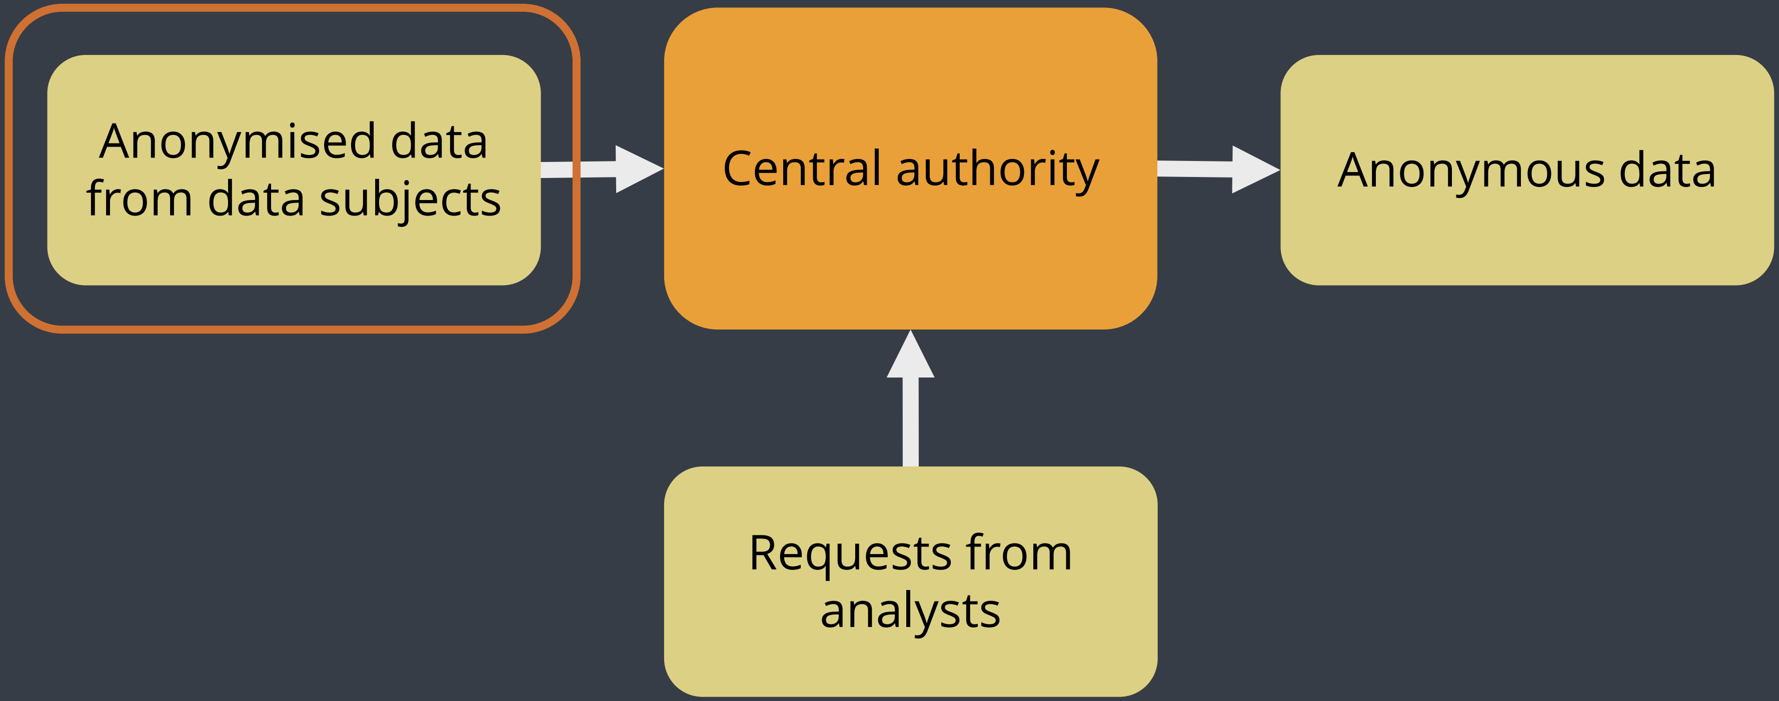

Anonymised data from data subjects

Central authority

Anonymous data

Requests from analysts

Google RAPPOR uses randomised response to prevent submissions being sensitive

Architecture

Untrusted central authority

Security definition

Differential privacy

Mechanism

Randomised response

Are you a member of the Communist Party?

1. Toss a fair coin
2. If heads, say “Yes”
3. If no, tell the truth

Gives **plausible deniability for individuals** but allows accurate aggregate statistics, and the uncertainty, to be calculated

RAPPOR extends this approach to arbitrary values, and shows that it fulfils differential privacy, subject to some caveats

Submissions are linked to individual identity though IP address, but are not sensitive even if linked to identity

RAPPOR offers strong security but is limited in terms of where it can be applied

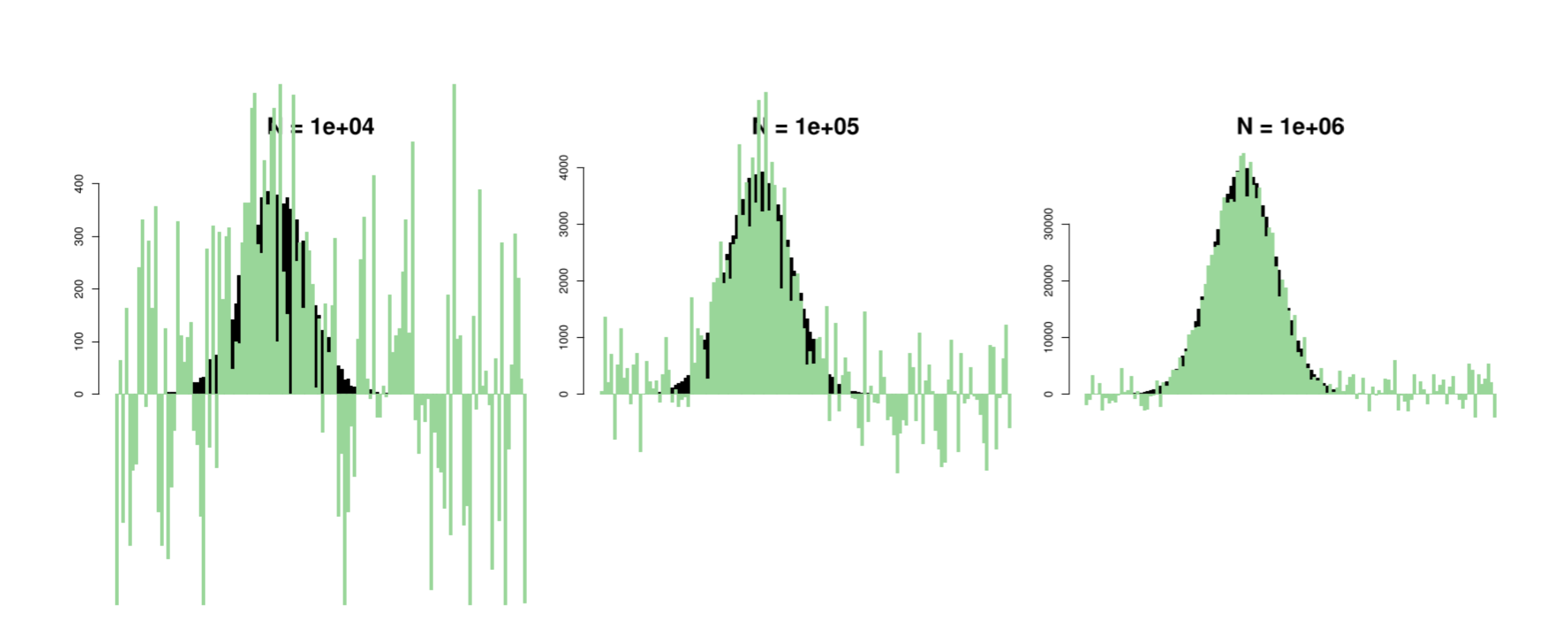

Homepage URL can be identified only if it is set by at least about 14,000 users which limits the type of organisations that can gain useful insights

Security definition assumes consecutive responses are uncorrelated

Anonymisation can be extended to include not only data but communications too

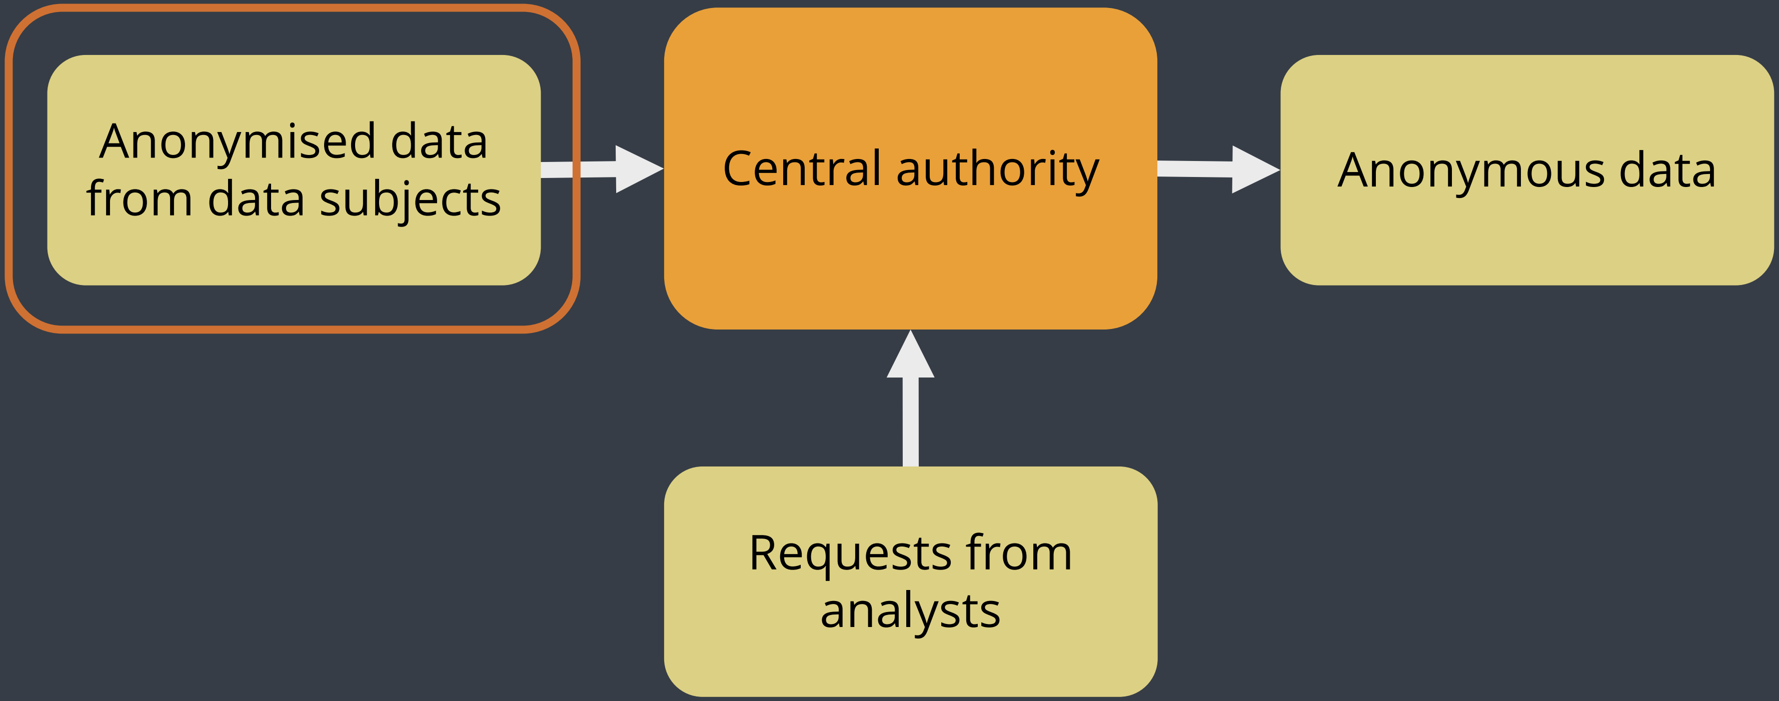

Anonymised data from data subjects

Central authority

Anonymous data

Requests from analysts

Anonymisation can be extended to include not only data but communications too

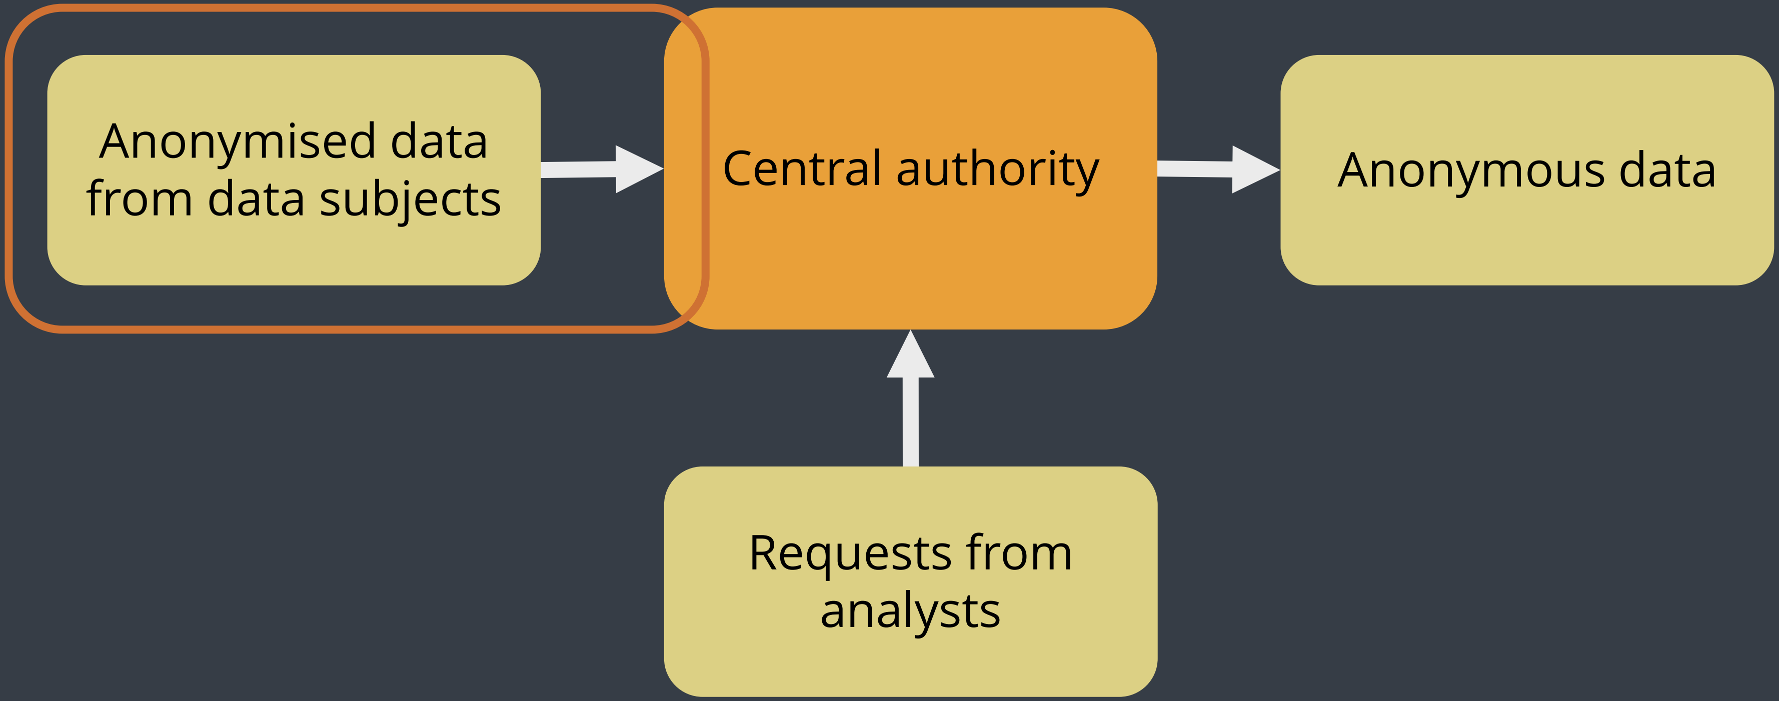

Anonymised data from data subjects

Central authority

Anonymous data

Requests from analysts

If data is only sensitive if linked to identity and communication is anonymous, the central authority is no longer trusted

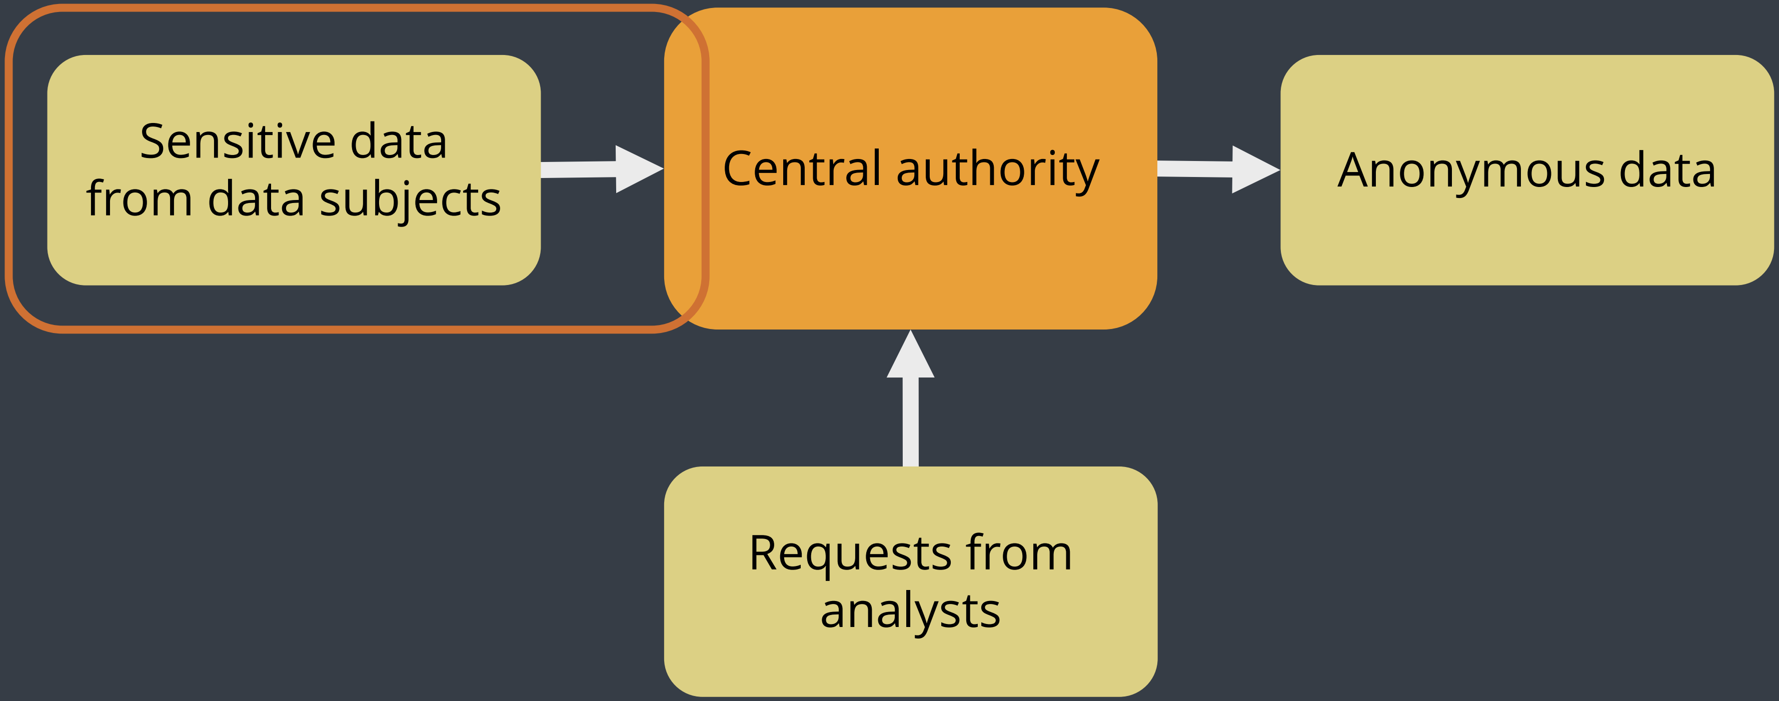

Sensitive data from data subjects

Central authority

Anonymous data

Requests from analysts

Short location traces can be anonymous if unlinked from other traces and identity

- An individual's location trace can be easily linked to an individual through auxiliary data
- Shown by Neustar Research using NYC Taxi dataset

+
−
Glen Rock
Fair Lawn
Paterson
Dumont
Yonkers
Eastchester
New Rochelle
Bayville
Hackensack
Glen Cove
Oyster Bay
Clifton
Verona
Bloomfield
295
25A
135
Union City
East Orange
495
Newark
Mineola
NASSAU
Hempstead
Elizabeth
Valley Stream
Freeport
NEW JERSEY
NEW YORK
Linden
878
Long Beach
W
Leaflet | © OpenStreetMap. Tiles courtesy of MapQuest

Splitting the location trace into small segments greatly reduces the risk of re-identification provided individual traces cannot be linked either by the trace or by the IP address the trace is submitted from

Intrusion detection logs can be anonymous if unlinked from other logs and identity

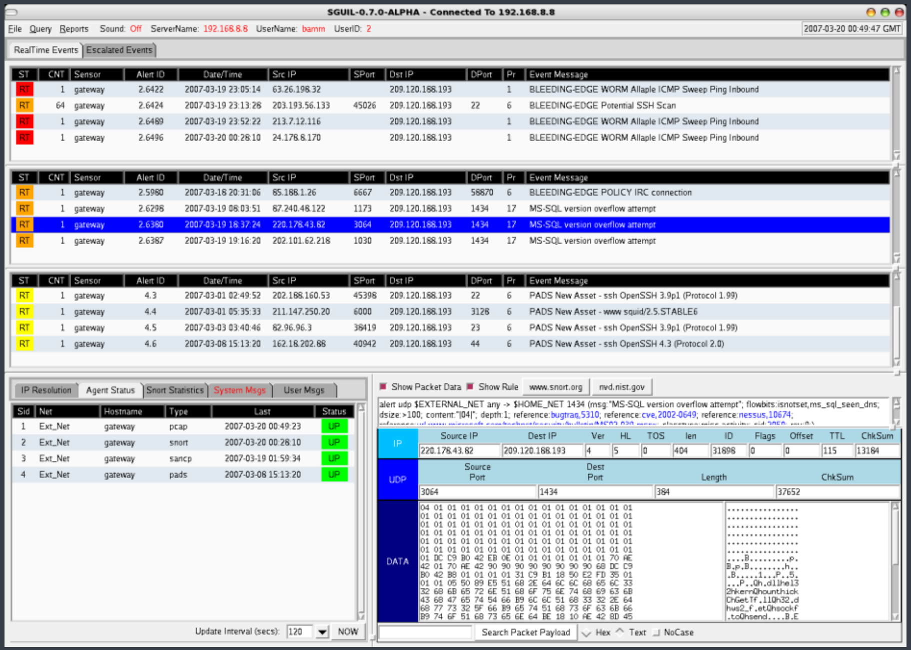

Sound: Off ServerName: 192.168.8.8 UserName: bamm UserID: 2

2007-03-20 00:49:47 GMT

ST	CNT	Sensor	Alert ID	Date/Time	Src IP	SPort	Dst IP	DPort	Pr	Event Message
RT	1	gateway	2.6422	2007-03-19 23:05:14	63.26.198.32		209.120.188.193		1	BLEEDING-EDGE WORM Allaple ICMP Sweep Ping Inbound
RT	64	gateway	2.6424	2007-03-19 23:13:28	203.193.56.133	45026	209.120.188.193	22	6	BLEEDING-EDGE Potential SSH Scan
RT	1	gateway	2.6489	2007-03-19 23:52:22	213.7.12.116		209.120.188.193		1	BLEEDING-EDGE WORM Allaple ICMP Sweep Ping Inbound
RT	1	gateway	2.6496	2007-03-20 00:28:10	24.178.8.170		209.120.188.193		1	BLEEDING-EDGE WORM Allaple ICMP Sweep Ping Inbound

ST	CNT	Sensor	Alert ID	Date/Time	Src IP	SPort	Dst IP	DPort	Pr	Event Message
RT	1	gateway	2.5960	2007-03-18 20:31:06	85.188.1.26	6667	209.120.188.193	58870	6	BLEEDING-EDGE POLICY IRC connection
RT	1	gateway	2.6298	2007-03-19 08:03:51	87.240.48.122	1173	209.120.188.193	1434	17	MS-SQL version overflow attempt
RT	1	gateway	2.6380	2007-03-19 18:37:24	220.178.43.82	3064	209.120.188.193	1434	17	MS-SQL version overflow attempt
RT	1	gateway	2.6387	2007-03-19 19:16:20	202.101.62.218	1030	209.120.188.193	1434	17	MS-SQL version overflow attempt

ST	CNT	Sensor	Alert ID	Date/Time	Src IP	SPort	Dst IP	DPort	Pr	Event Message
RT	1	gateway	4.3	2007-03-01 02:49:52	202.188.160.53	45398	209.120.188.193	22	6	PADS New Asset - ssh OpenSSH 3.9p1 (Protocol 1.99)
RT	1	gateway	4.4	2007-03-01 05:35:33	211.147.250.20	6000	209.120.188.193	3128	6	PADS New Asset - www squid/2.5.STABLE6
RT	1	gateway	4.5	2007-03-03 03:40:46	82.96.96.3	38419	209.120.188.193	23	6	PADS New Asset - ssh OpenSSH 3.9p1 (Protocol 1.99)
RT	1	gateway	4.6	2007-03-08 15:13:20	162.18.202.88	40942	209.120.188.193	44	6	PADS New Asset - ssh OpenSSH 4.3 (Protocol 2.0)

Sid	Net	Hostname	Type	Last	Status
1	Ext_Net	gateway	pcap	2007-03-20 00:49:23	UP
2	Ext_Net	gateway	snort	2007-03-20 00:28:10	UP
3	Ext_Net	gateway	sancp	2007-03-19 01:59:34	UP
4	Ext_Net	gateway	pads	2007-03-08 15:13:20	UP

Update Interval (secs): 120

```
alert udp $EXTERNAL_NET any -> $HOME_NET 1434 (msg:"MS-SQL version overflow attempt"; flowbits:isnotset,ms_sql_seen_dns;
dsize:>100; content:"|04|"; depth:1; reference:bugtraq,5310; reference:cve,2002-0649; reference:nessus,10674;
```

IP	Source IP	Dest IP	Ver	HL	TOS	len	ID	Flags	Offset	TTL	ChkSum
	220.178.43.82	209.120.188.193	4	5	0	404	31898	0	0	115	13184

UDP	Source Port	Dest Port	Length	ChkSum
	3064	1434	384	37652

DATA

```
04 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 01 01 01 01 01 01 01 01 01 01 01 01 01 01 01    ................
01 DC C9 B0 42 EB 0E 01 01 01 01 01 01 01 70 AE    ....B.........p.
42 01 70 AE 42 90 90 90 90 90 90 90 90 68 DC C9    B.p.B........h..
B0 42 B8 01 01 01 01 31 C9 B1 18 50 E2 FD 35 01    .B.....1...P..5.
01 01 05 50 89 E5 51 68 2E 64 6C 6C 68 65 6C 33    ...P..Qh.dllhel3
32 68 6B 65 72 6E 51 68 6F 75 6E 74 68 69 63 6B    2hkernQhounthick
43 68 47 65 74 54 66 B9 6C 6C 51 68 33 32 2E 64    ChGetTf.llQh32.d
68 77 73 32 5F 66 B9 65 74 51 68 73 6F 63 6B 66    hws2_f.etQhsockf
B9 74 6F 51 68 73 65 6E 64 BE 18 10 AE 42 8D 45    .toQhsend....B.E
```

Organisations may wish to share information on attacks they are subject to but still hide their identity

Real example of a consortium of large companies, competing with each other but subject to same threats

Techniques like prefix-preserving anonymization can hide details from logs but still need to submit data without disclosing identity

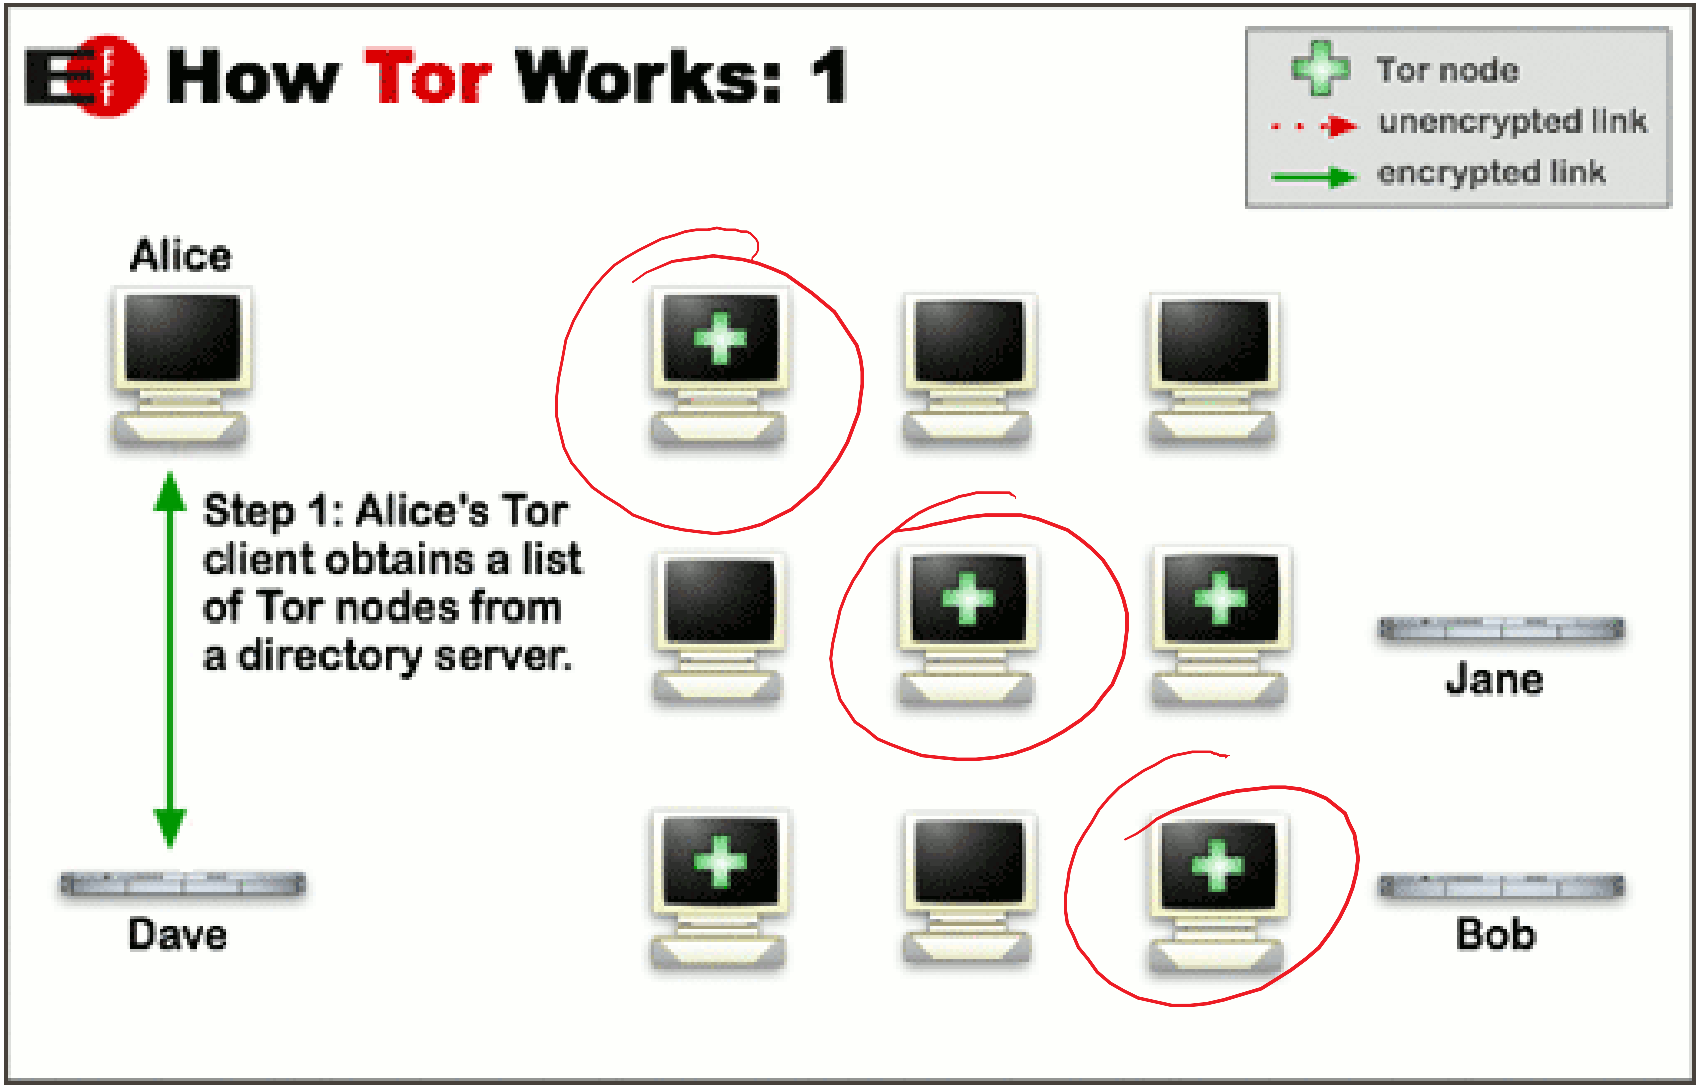
How Tor Works: 1
Tor node
unencrypted link
encrypted link
Alice
Step 1: Alice's Tor client obtains a list of Tor nodes from a directory server.
Jane
Dave
Bob

Network nodes are run by volunteers, and changes to structure are made by consensus decision of 8 directory authority operators

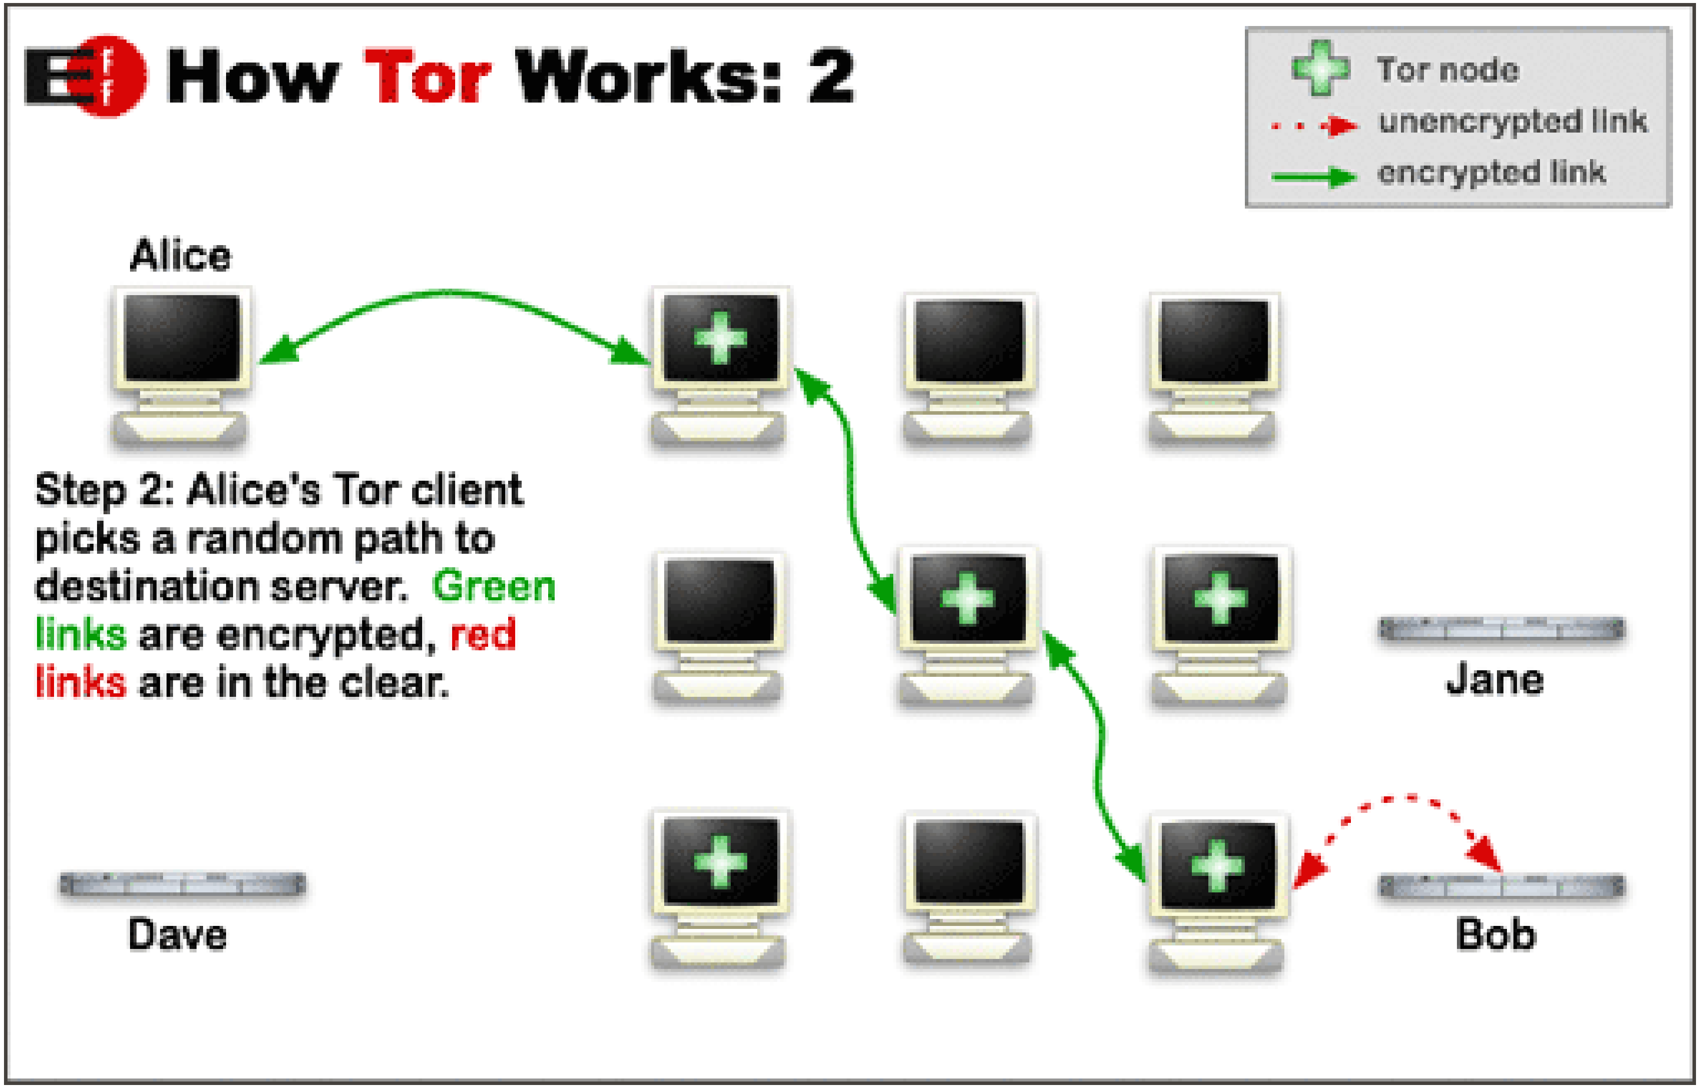

How Tor Works: 2

Step 2: Alice's Tor client picks a random path to destination server. Green links are encrypted, red links are in the clear.

Each node knows the prior step in the circuit, and the next step, but unless there is collusion nobody can link source to destination

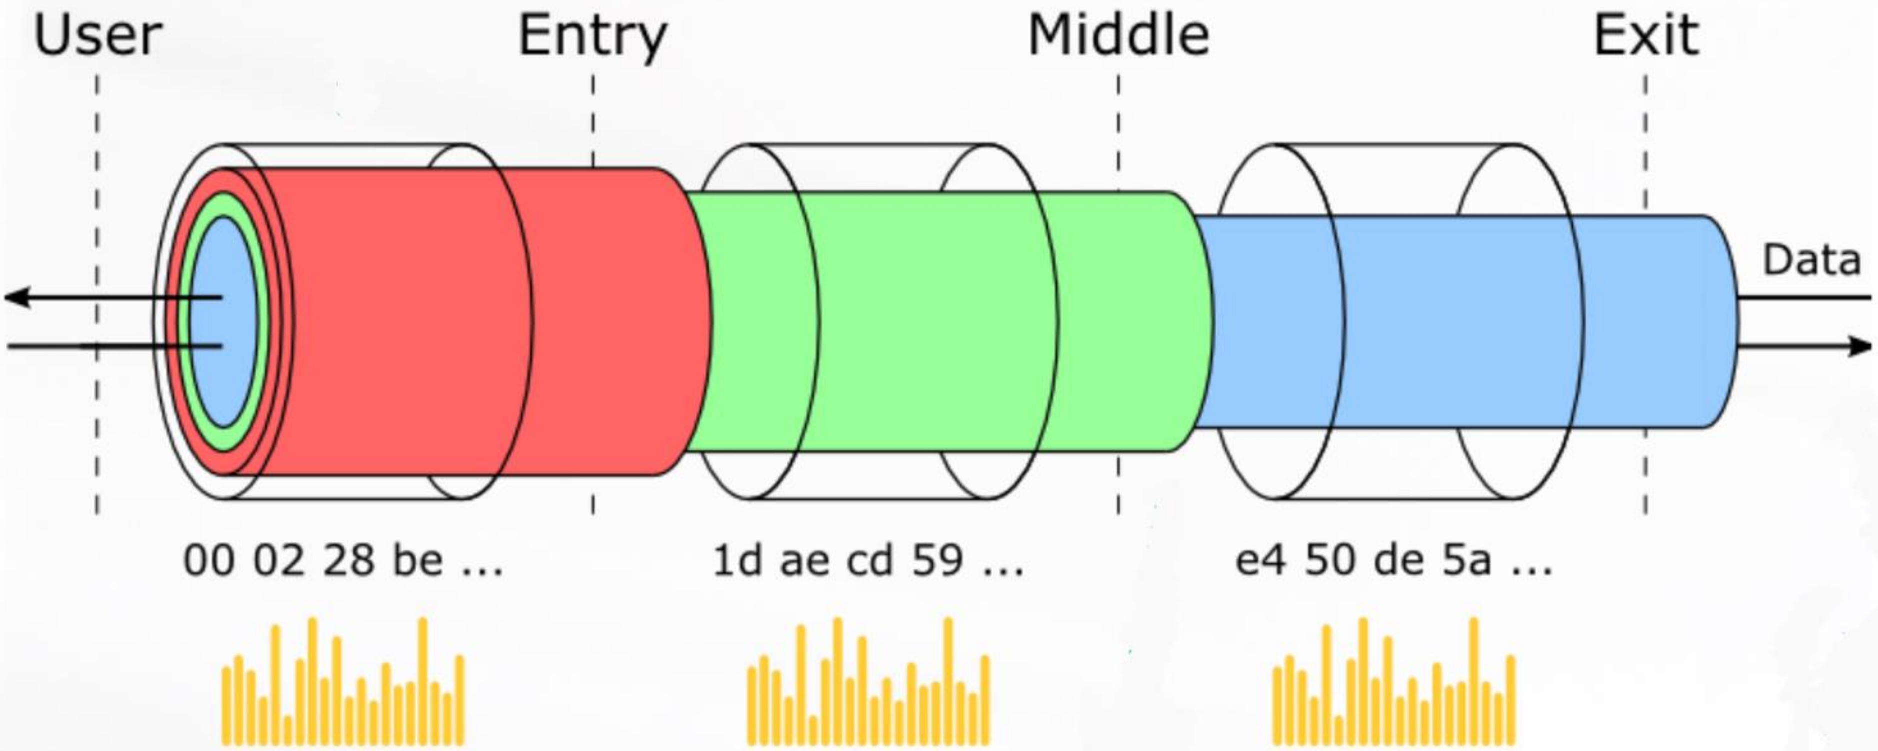

Encryption prevents data flowing into a node from being linked to data flowing out, preserving unlinkability property

Tor and other common anonymity systems do not hide traffic patterns

There are useful alternatives between the extremes of fully centralised and full anonymisation by data subject

- Fully centralised approaches are versatile but have disadvantages
  - Risk of data compromise, expensive to set up and maintain security
  - Public lacks confidence in institutions
  - Centralisation of power lends itself to authoritarian tendencies
- Local anonymisation such as RAPPOR has limitations
  - Requires large amounts of data to get meaningful results
  - Care needed to preserve security if there is correlated data
- Data can be anonymised such that it is not sensitive in itself but does not offer plausible deniability either
  - In which case, submitting data over an anonymous communication system such as Tor can protect privacy without having to trust a central authority to anonymise data