That’s rather inappropriate, dear:

The challenges of determining ‘appropriate’ measures for personal data deletion

PRESENTER

Cat Easdon

Senior Privacy Engineer

Who am I?

- Senior Privacy Engineer at Dynatrace
- ⛰-obsessed Brit based in Innsbruck
- Previously: hacking CPUs at TU Graz
- Still have one foot in academia

Outline

- ‘Reasonableness’ and ‘Appropriateness’
- Case study 1: remediation
  - The unexpected challenges of cleaning up personal data in source control
- Case study 2: retention
  - How far should we go when deleting accounts?

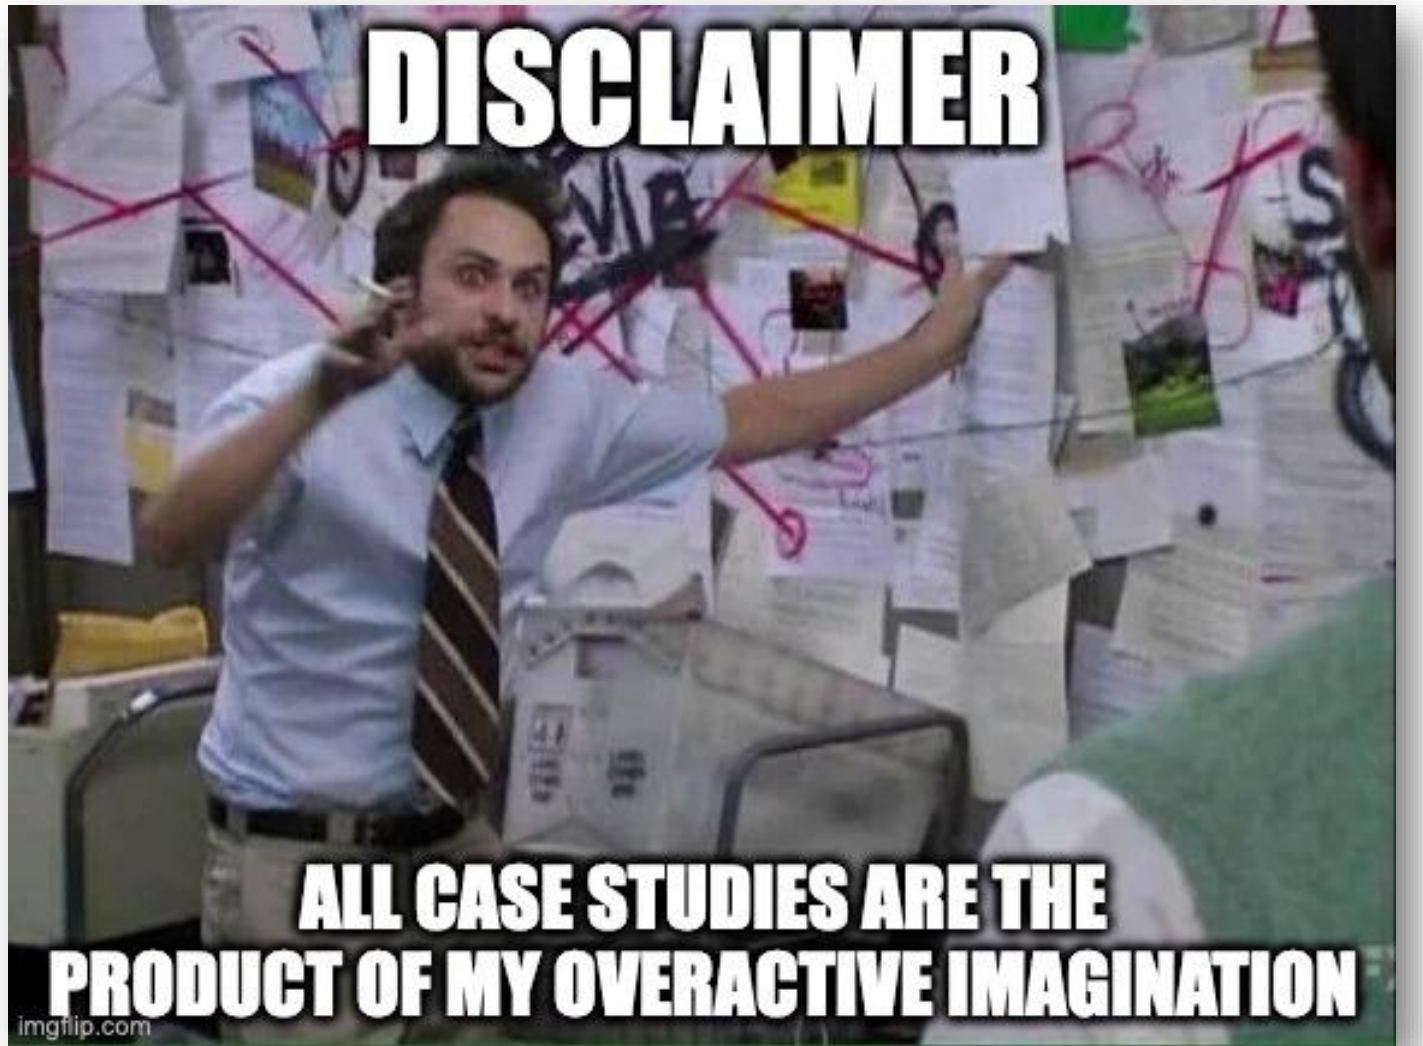
DISCLAIMER
ALL CASE STUDIES ARE THE
PRODUCT OF MY OVERACTIVE IMAGINATION
imgflip.com

‘Reasonableness’ and ‘Appropriateness’

‘Reasonable’ and ‘appropriate’ measures: GDPR

‘The protection of the rights and freedoms of natural persons with regard to the processing of personal data require that **appropriate technical and organisational measures** be taken to ensure that the requirements of this Regulation are met…the controller should **adopt internal policies and implement measures** which meet in particular the principles of data protection by design and data protection by default.’

‘Reasonable’ and ‘appropriate’ measures: CCPA and POPIA

- **CCPA:** ‘implement reasonable security measures’ to protect consumers’ personal information (incl. in transit) and privacy rights request records and to detect fraudulent identity verification
- **POPIA:** appropriate, reasonable technical and organisational security measures
  - Establish, maintain, verify, and update safeguards to prevent loss of, damage to, unauthorized destruction of, unlawful access to, or unlawful processing of personal information
  - Having due regard to generally accepted information security practices and procedures

But what is ‘reasonable’?

- CCPA: This is ‘a fact-specific determination’
- ‘It would be too limiting to prescribe reasonable security measures’
- ‘Consult with an attorney who is aware of all pertinent facts and relevant compliance concerns’

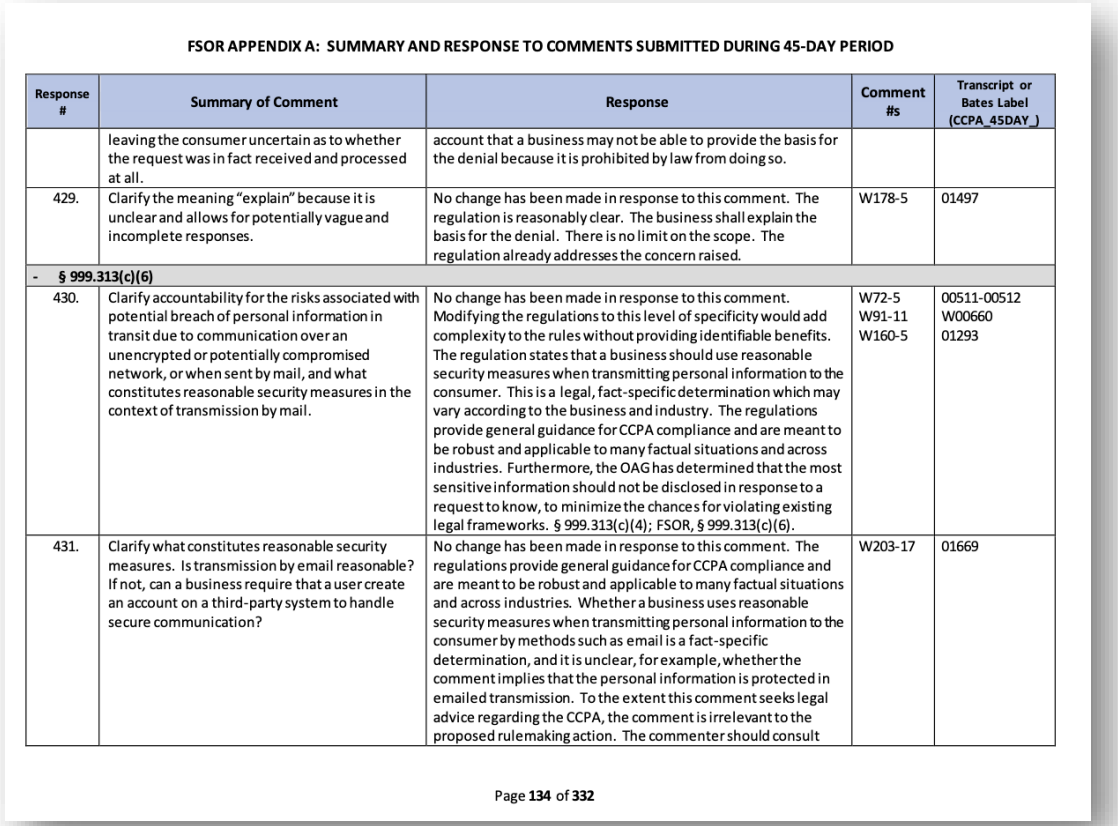

FSOR APPENDIX A: SUMMARY AND RESPONSE TO COMMENTS SUBMITTED DURING 45-DAY PERIOD

Response #	Summary of Comment	Response	Comment #s	Transcript or Bates Label (CCPA_45DAY_)
	leaving the consumer uncertain as to whether the request was in fact received and processed at all.	account that a business may not be able to provide the basis for the denial because it is prohibited by law from doing so.		
429.	Clarify the meaning “explain” because it is unclear and allows for potentially vague and incomplete responses.	No change has been made in response to this comment. The regulation is reasonably clear. The business shall explain the basis for the denial. There is no limit on the scope. The regulation already addresses the concern raised.	W178-5	01497
- § 999.313(c)(6)				
430.	Clarify accountability for the risks associated with potential breach of personal information in transit due to communication over an unencrypted or potentially compromised network, or when sent by mail, and what constitutes reasonable security measures in the context of transmission by mail.	No change has been made in response to this comment. Modifying the regulations to this level of specificity would add complexity to the rules without providing identifiable benefits. The regulation states that a business should use reasonable security measures when transmitting personal information to the consumer. This is a legal, fact-specific determination which may vary according to the business and industry. The regulations provide general guidance for CCPA compliance and are meant to be robust and applicable to many factual situations and across industries. Furthermore, the OAG has determined that the most sensitive information should not be disclosed in response to a request to know, to minimize the chances for violating existing legal frameworks. § 999.313(c)(4); FSOR, § 999.313(c)(6).	W72-5 W91-11 W160-5	00511-00512 W00660 01293
431.	Clarify what constitutes reasonable security measures. Is transmission by email reasonable? If not, can a business require that a user create an account on a third-party system to handle secure communication?	No change has been made in response to this comment. The regulations provide general guidance for CCPA compliance and are meant to be robust and applicable to many factual situations and across industries. Whether a business uses reasonable security measures when transmitting personal information to the consumer by methods such as email is a fact-specific determination, and it is unclear, for example, whether the comment implies that the personal information is protected in emailed transmission. To the extent this comment seeks legal advice regarding the CCPA, the comment is irrelevant to the proposed rulemaking action. The commenter should consult	W203-17	01669

Page **134** of **332**

But what is ‘reasonable’?

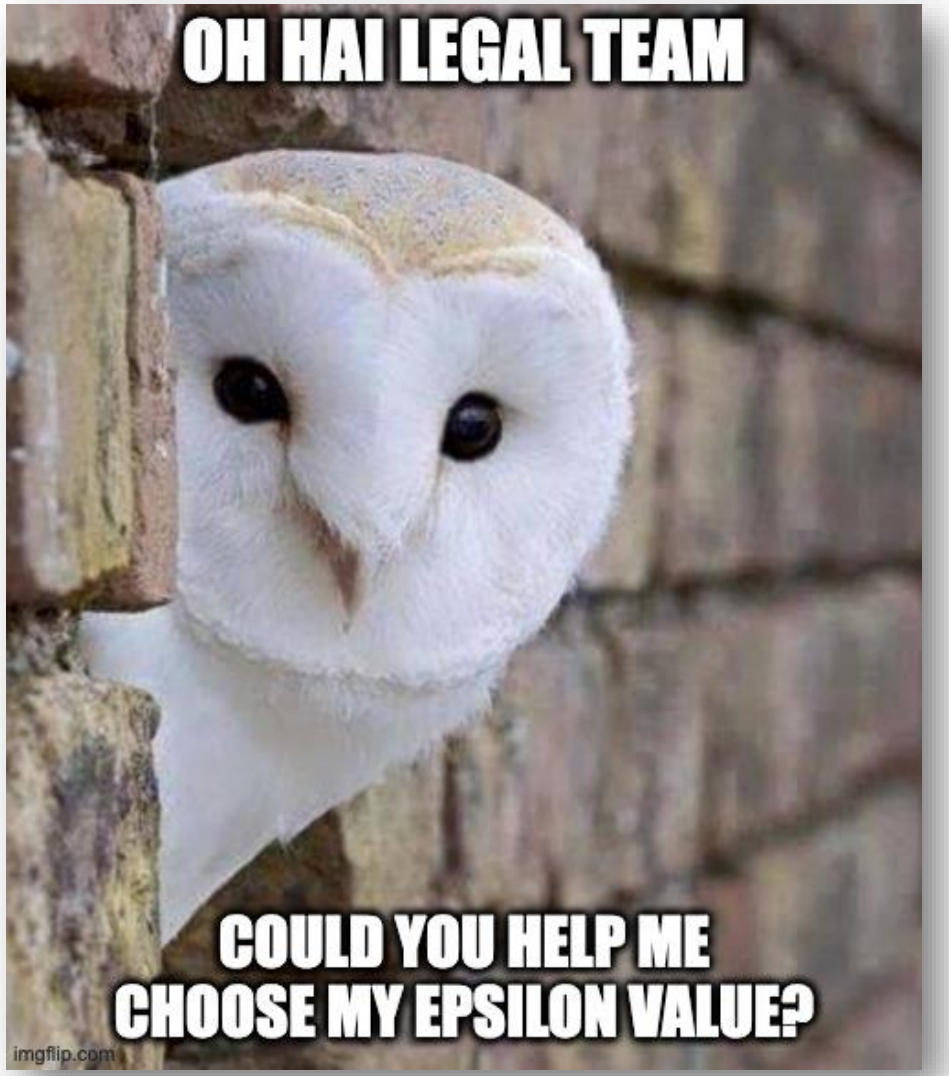

Thankfully, we do have some precedents…

- Center for Internet Security's 20 Critical Security Controls (identified as a baseline in the 2016 California Data Breach Report)
- CPPA's proposed cybersecurity audit regulation lists basic controls (including masking, retention periods, and data flow mapping)
- Sector-specific standards
- General standards (ISO, NIST…)
- Guidance from regulatory authorities
- Enforcement cases
- Industry best practices

Reasonable security and privacy according to the US FTC

- **Encrypt data**
- Mitigate known vulnerabilities
- Enforce good credential practices
- Use MFA
- Monitor and control network access
- **Maintain a written security program** (includes a data retention program)
- **Maintain a vulnerability disclosure program**
- Patch systems
- **Perform testing and auditing**
- **Minimize data retention and access**
- **Oversee service providers**
- **Train employees and personnel**

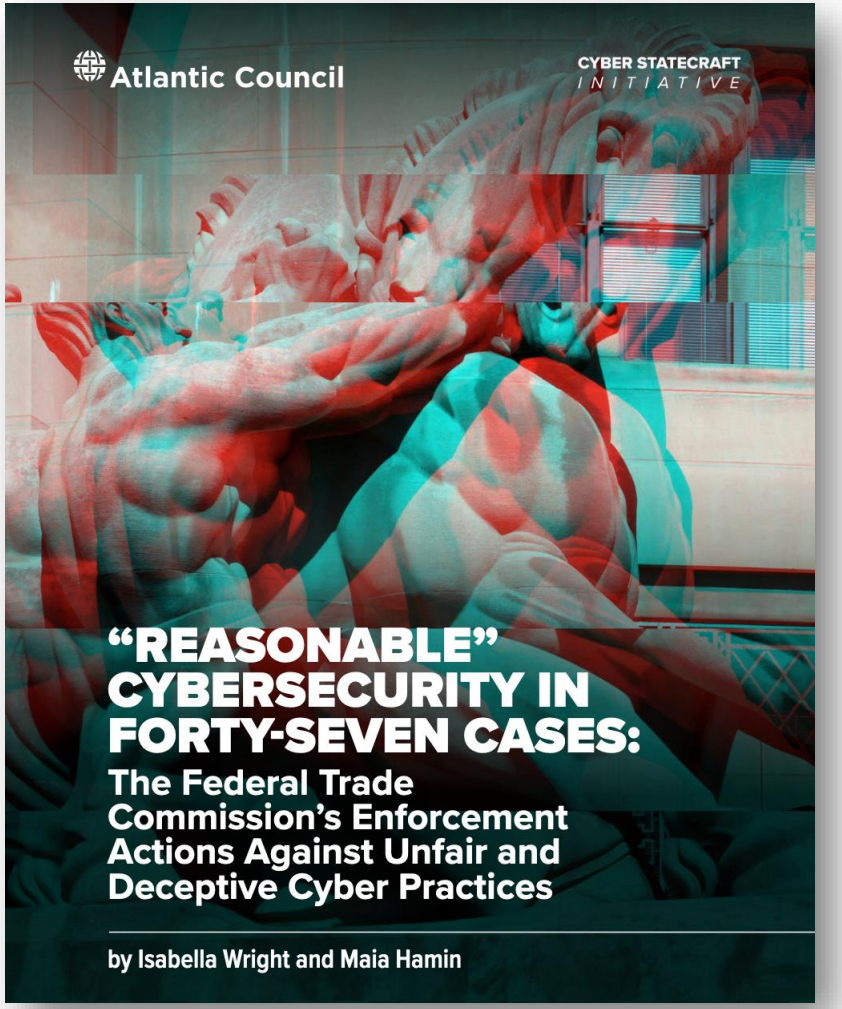

GDPR: appropriate measures

Interpretation based on Article 25, Recital 78, Article 32, Recital 83, focusing on the SDLC

- Pseudonymize (as soon as possible) and encrypt personal data
- Privacy as the default setting
- Purpose limitation by default
  - Prevent individuals from repurposing personal data
  - Robust audit logging, access control with least privilege
  - **Minimize volume of data and extent of processing, establish retention periods**
- Security controls for CIA and resilience, incl. disaster recovery process and regular testing
- Be transparent with data subjects about data processing and allow them to monitor it (where appropriate)
- Document all measures and policies; train employees; clearly assign responsibilities

GDPR: appropriate retention

> 2. The controller shall implement appropriate technical and organisational measures for ensuring that, by default, only personal data which are necessary for each specific purpose of the processing are processed. That obligation applies to the amount of personal data collected, the extent of their processing, the period of their storage and their accessibility. In particular, such measures shall ensure that by default personal data are not made accessible without the individual's intervention to an indefinite number of natural persons.

'Appropriate' taking into account:

- State of the art
- Cost of implementation
- Nature, scope, context, purposes of processing
- Risks to rights and freedoms of data subjects

But *how* exactly should these be taken into account?

(Does anyone have a formula?...)

Case Study 1: Remediation

Background: version control with git

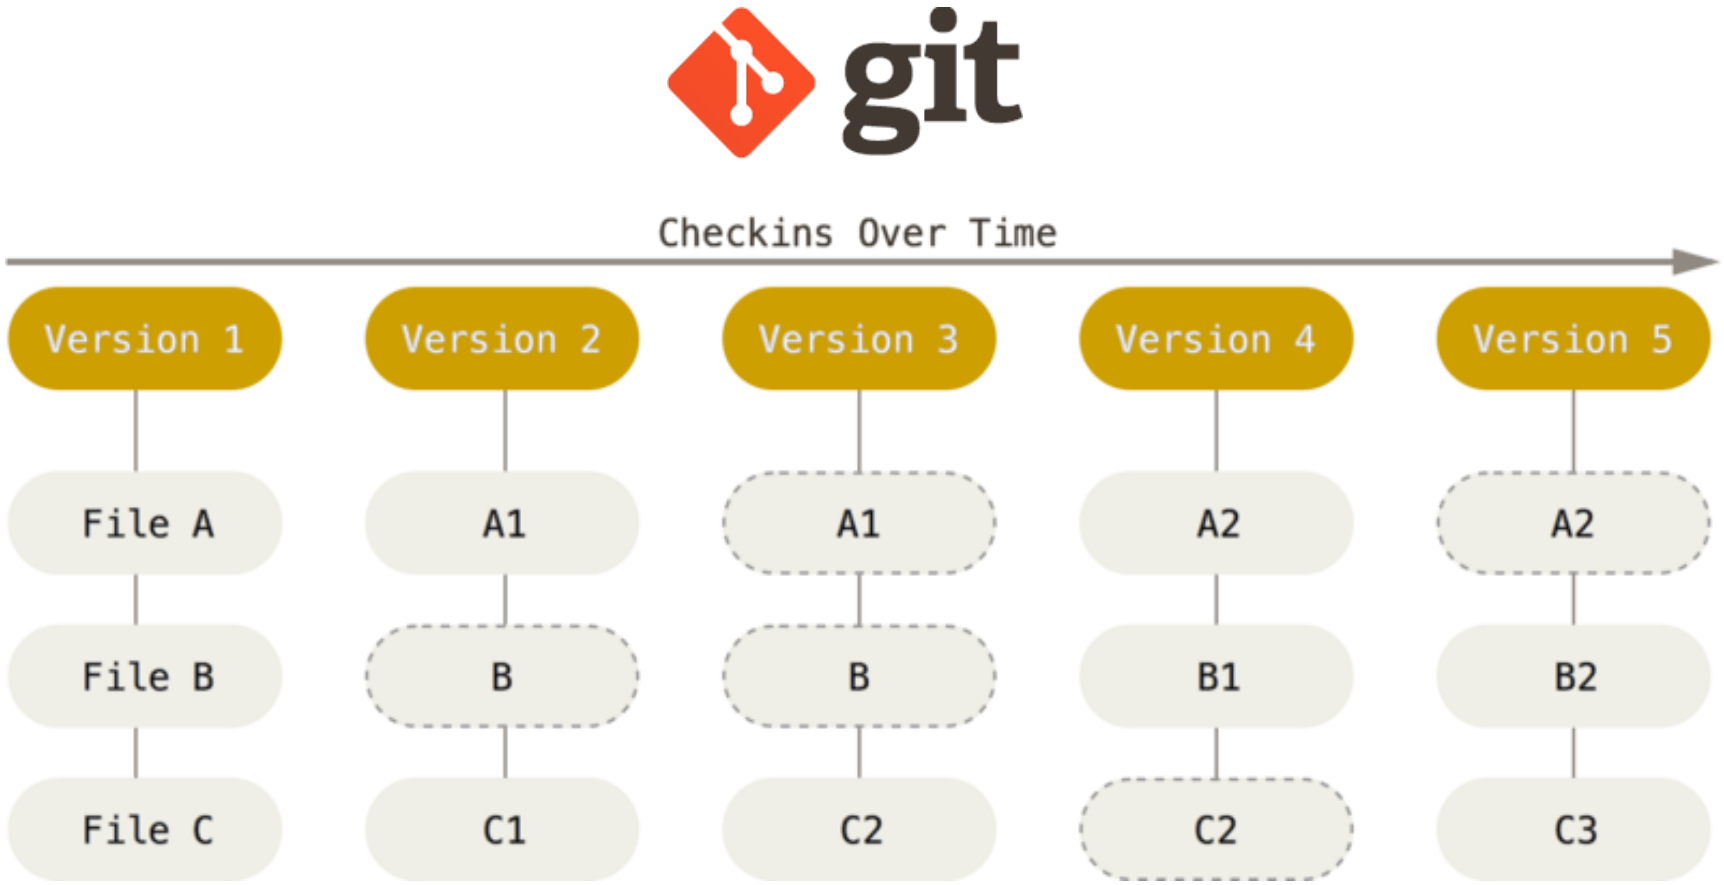

Background: version control with git

97c2c45253b	NOISSUE: Update dependency react-router-dom from 6.22.3 to 6.23.0	06 May 2024
abfb8f6bc57	NOISSUE: Update dependency @babel/preset-env from 7.24.4 to 7.24.5	06 May 2024
7b231b3b99c	NOISSUE: Update dependency styled-components from 6.1.8 to 6.1.9	06 May 2024
c5e0b514168	Pull request #308: NOISSUE: Update dependency i18n-iso-countries from 7.11.0 to 7.11.1 Sq	06 May 2024
c6bf3fd6ff4	NOISSUE: Update Dynatrace SDK	28 Apr 2024
4714ccfc330	NOISSUE: Update dependency eslint-plugin-no-secrets from 0.8.9 to 0.9.1	23 Apr 2024
0ab12416134	NOISSUE: Update dependency @lwc/eslint-plugin-lwc from 1.7.2 to 1.8.1	29 Apr 2024
14127bab859	NOISSUE: Update Dynatrace ecosystem	05 May 2024
36171f33331	Pull request #284: NOISSUE: Update react monorepo Squashed commit of the following: co	05 May 2024
07964a2cafd	Pull request #301: NOISSUE: Update storybook monorepo from 7.6.17 to 7.6.19 Squashed co	05 May 2024
6ee3262fc0c	[redacted]: Accessibility improvements for forms [redacted] Accessibility improvements for	03 May 2024
9e648e63d9d	[redacted]: Add aria-label to DefaultPolicyIndicator. [redacted]: Add aria-label to DefaultPoli	30 Apr 2024

Version control ❤️ a spreadsheet full of personal data

Challenges: compliance requirements and deletion itself

System is used for customer billing

- Enter stage left: Sarbanes-Oxley Act (SOX)
- History should be **immutable** for change management controls compliance

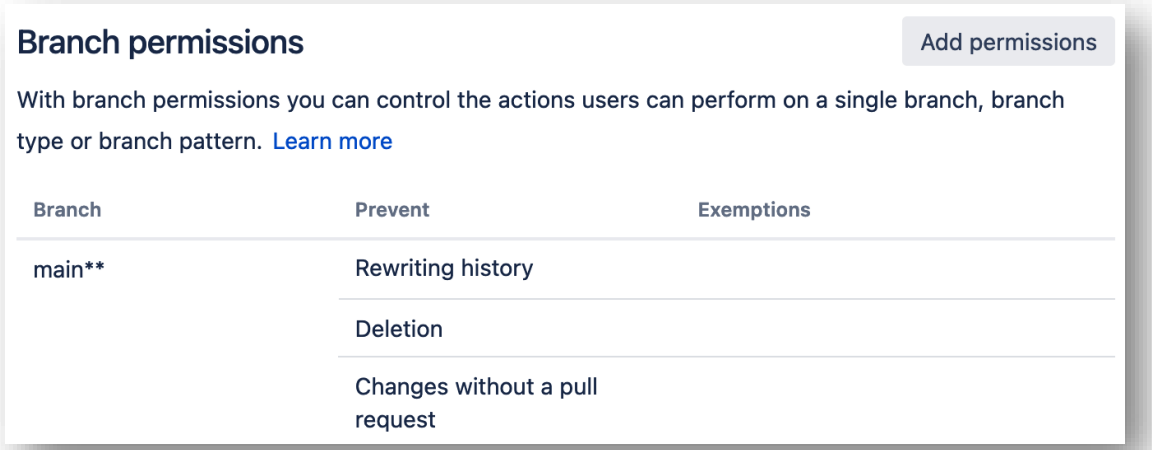

Git doesn't natively support erasing commits

- Research needed to identify tooling and assess risk of data loss

Challenge: avoiding accidental releases

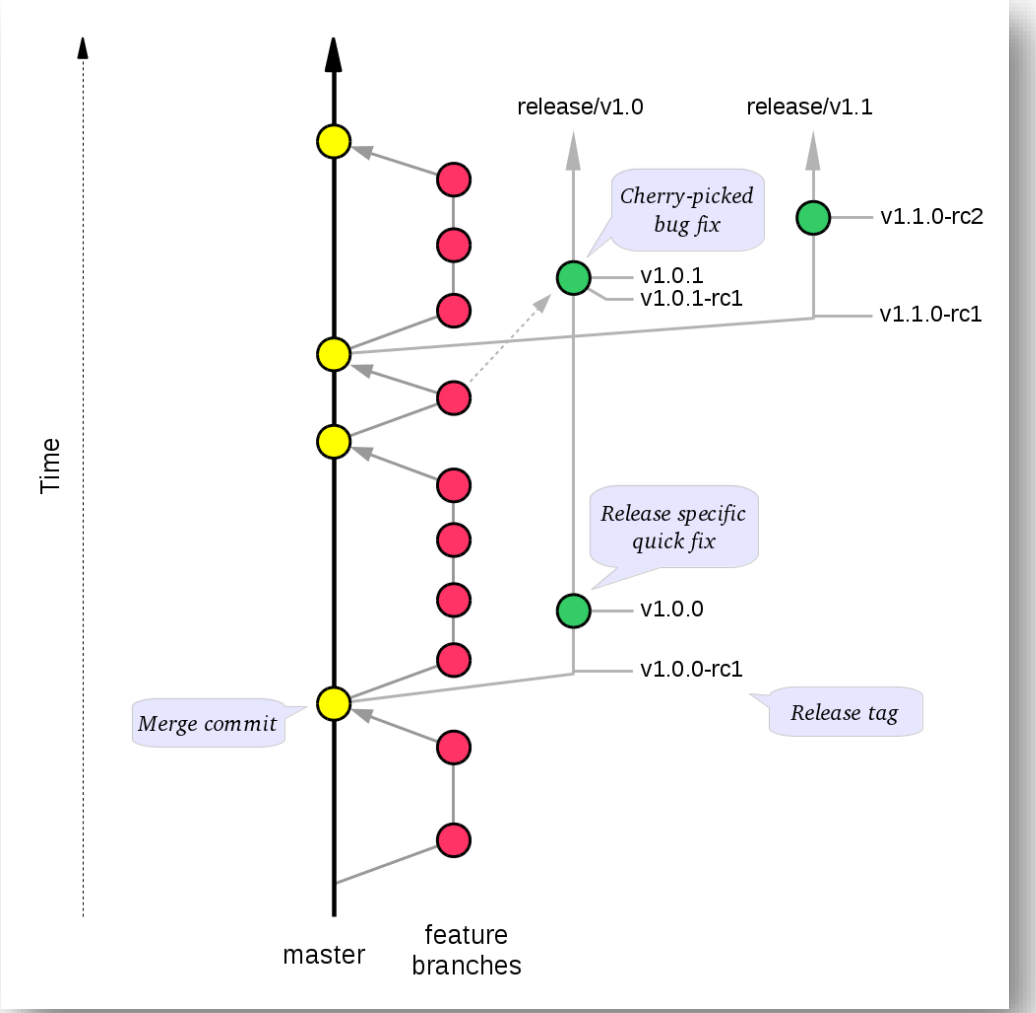

Challenge: how do we eradicate the orphaned commits?

Further deletion requires manually triggering garbage collection

- May cause data loss
- Requires weekend overtime for infra team
- First attempt fails, long support case with Atlassian
- Interim mitigation: block all access to (internally-hosted) Bitbucket URL for the commit

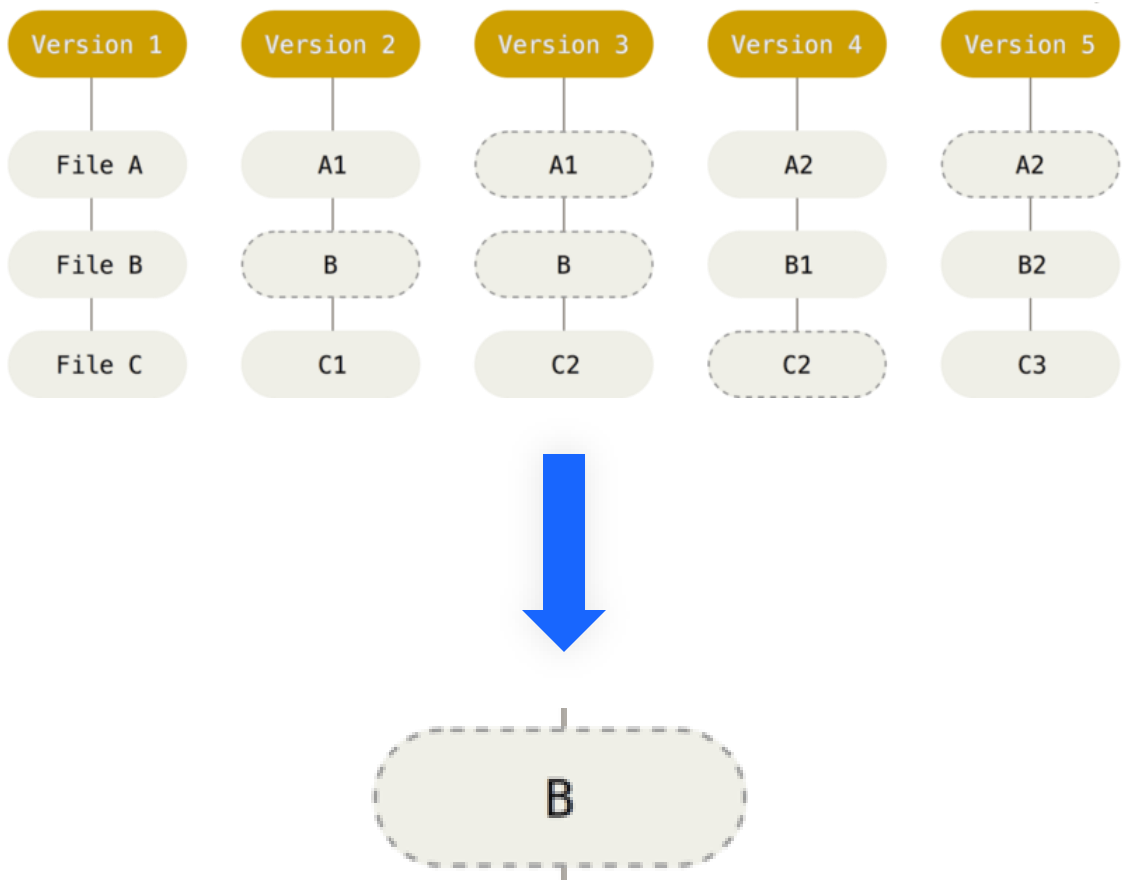

“Taking into account…the cost of implementation”

Dev time for repo and pipeline research and deletion: 10 hours

Infra team time to delete orphaned commit: 5 hours, incl. weekend overtime

Privacy consulting and negotiation: 3 hours

Setting up dev environment again: 30 minutes x 50 devs + 1hr communication = 26 hours

Total: 44 hours to delete a spreadsheet ☹

+ tooling research (one-time only): 2hrs

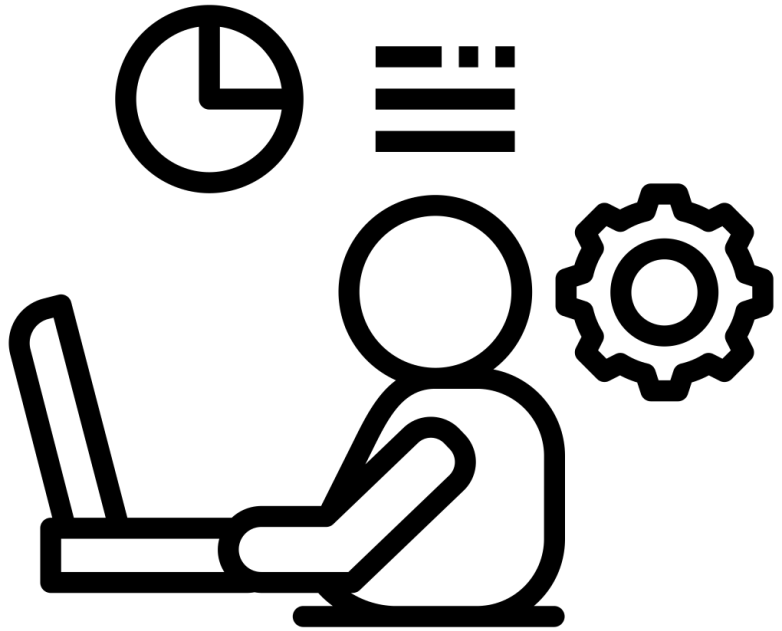

Risks to rights and freedoms vs. cost of implementation

Risks:

- Original pull request deleted
- Not searchable once removed from all branches
- Access to repo is limited

```
64ae98dfeeaf242d3515067fd2f00afbf0708d35
```

Cost:

- Increases drastically for larger repos (with 500 devs paid €30/hr: 251hrs = €7530 just for dev environment setup)
- Process costs goodwill within the company, setting back roadmaps
- Devs argue: if we can't be certain the data's eradicated, it's still illegal, so what's the point?

Case Study 2: Retention

Account deletion

- Deletion request or end of a free trial: time to clean up user and account data
- Clearly, we need to establish deletion mechanisms and retention policies in each data store
- Requirement: no unlimited free trials
- Conflicting compliance requirements, e.g. audit logging and billing records
  - Again: “if we retain *something*, we might as well retain *everything*…”

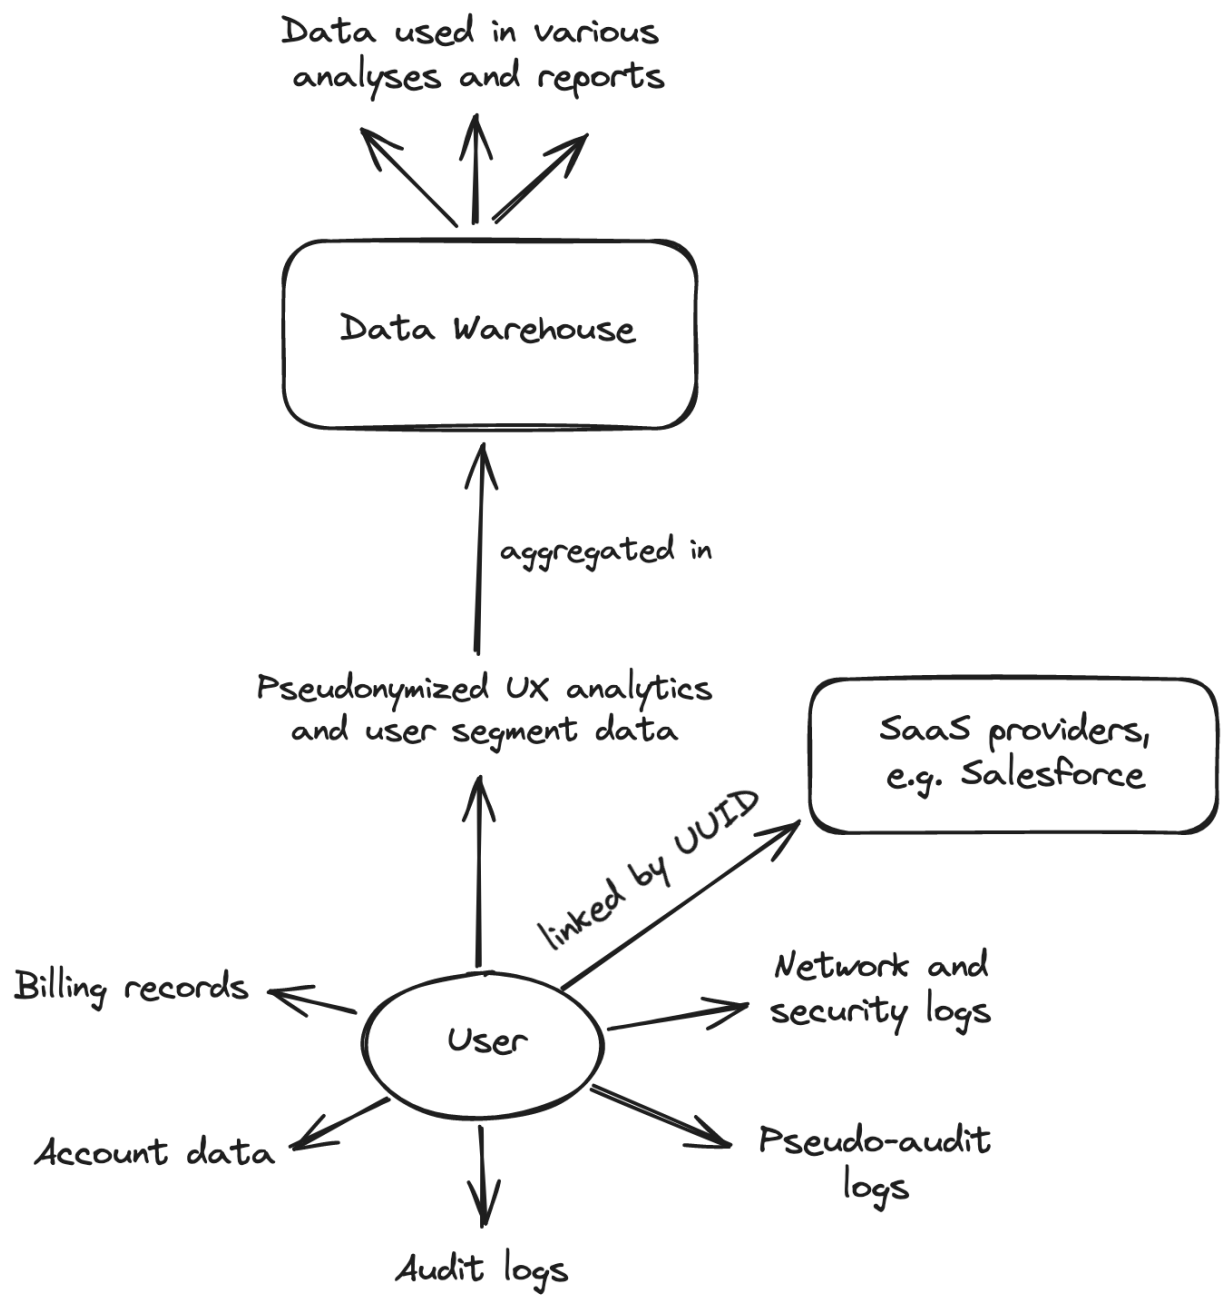

Magically delete everything!

Deletion requests

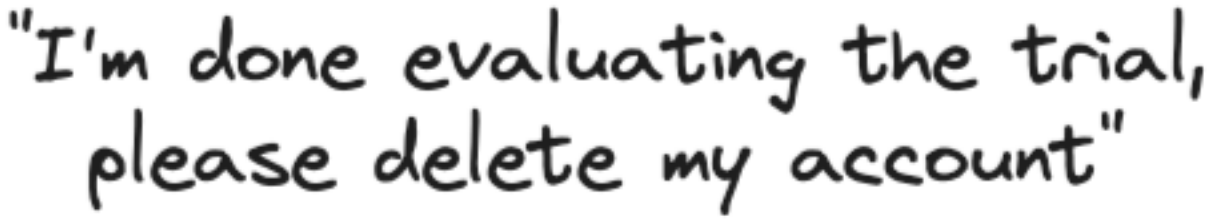

- Is this even a data subject rights request? Will the user understand if we ask them to confirm that?
- Sales would like to contact this user to find out what we could do better
- UX, Marketing, and Product teams would like data about this user's experience and their user segment
- Recital 47: *"The interests and fundamental rights of the data subject could…override the interest of the data controller where personal data are processed in circumstances where data subjects* ***do not reasonably expect*** *further processing."*

What is the risk to the data subject?

LINDDUN

Privacy Threat
Linkability
Identifiability
Non-Repudiation
Detectability
Data Disclosure
Unawareness, Unintervenability
Non-Compliance

Privacy Harms
Physical Harms
Economic Harms
Reputational Harms
Autonomy Harms
Psychological Harms
Discrimination Harms
Relationship Harms
Harm to the Business

Autonomy Harms
Coercion
Manipulation
Failure to Inform
Thwarted Expectations
Lack of Control
Chilling Effects

Privacy Harms, Citron & Solove, Boston University Law Review (2022)

Key takeaways

1. For a small number of security and privacy controls, what's 'appropriate' or 'reasonable' is surprisingly well-defined.
   - Advocate for these controls in your organization, if they're not in place already

2. For everything else, a risk-based approach is crucial, not only for *making* privacy decisions, but also for *negotiating* and *communicating* privacy measures within your organization.
   - We can't afford to abandon the risk-based approach!
   - Threat modeling can help us make these tough judgment calls and persuade others
   - Anticipate and counter the argument *"But we're bound to have some trace of personal data left somewhere…We'll still be breaking the law, so what's the point even trying?".*

Thanks for listening! Any questions?

Copyright notice

- **Dynatrace content and branding:** © 2024 Dynatrace LLC
- **Third-party images, text, and videos:** see links for attribution
- **Unattributed images:** generated with DALL·E 3 or used under license from the Noun Project
- **Diagrams:** created with Excalidraw
- **All other content:** original work by the author, may be reused with attribution