Smals & Postgres

“One flew over the cuckoo’s nest”
aka
Some Insights

Dirk Deridder

Smals
ICT for society

Smals
ICT for society

- Belgian **In-house** ICT Service Provider/Integrator
- For and by the **Public Sector**
- **Nonprofit** organization (VZW/ASBL)
- **2,166** collaborators
- Operates as a **shared service center** for **319** government organizations
  - Infrastructure management
  - Software development
  - Data management
  - Security services
  - Recruitment services …
- Main role: **enabler** for government ICT
- Established in **1939**

More than 80 years of data & information management

ACTIVITEITENVERSLAG
2023

Smals
ICT for society

https://www.smals.be/nl/content/activiteitenverslag-2023

Pre-Postgres situation…

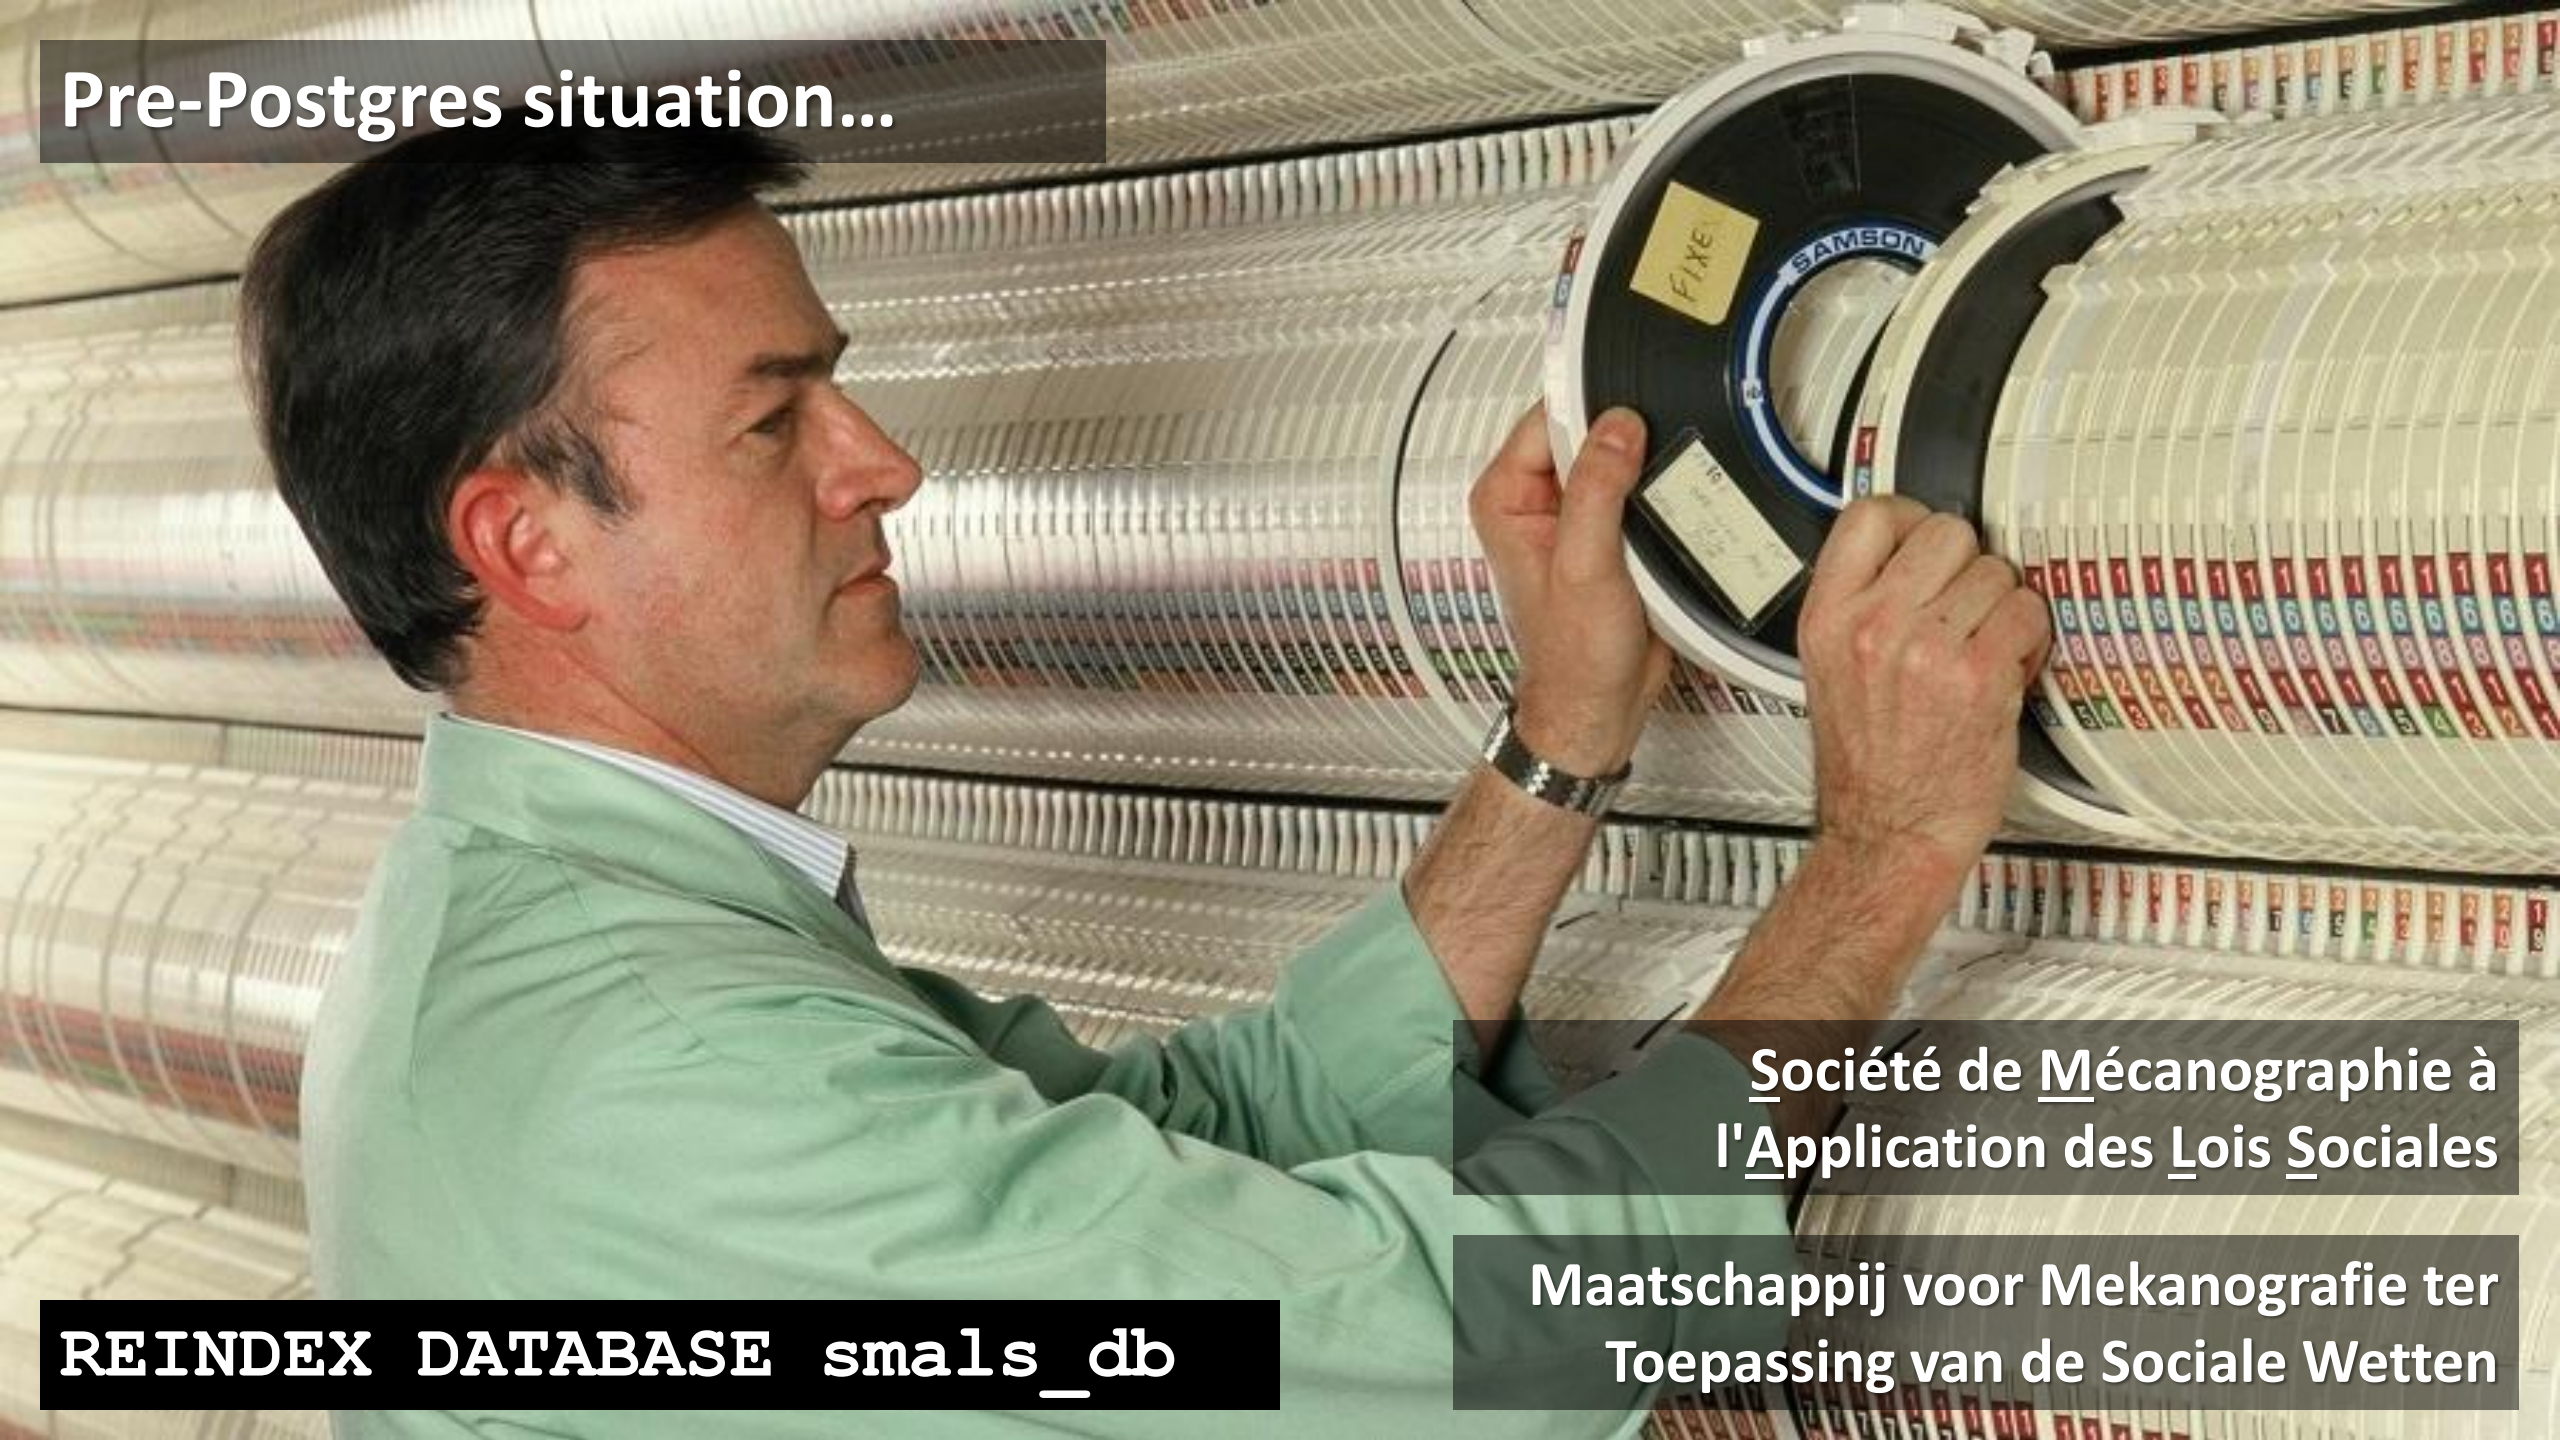

Société de Mécanographie à l'Application des Lois Sociales

Maatschappij voor Mekanografie ter Toepassing van de Sociale Wetten

```
REINDEX DATABASE smals_db
```


Smals
ICT for society
WIE WE ZIJN
WAT WE DOEN
NIEUWS
REALISATIES
JOBS
LEVERANCIERS
CONTACT
NL
Post-Postgres situation…
16 YEARS
top EMPLOYER
BELGIË
BELGIQUE
BELGIUM
2024
CERTIFIED EXCELLENCE IN EMPLOYEE CONDITIONS
Hey Siri, please improve the performance of the way we handle data and information at Smals.
Make sure to avoid any business impact, but to be sure inform all customers upfront. Afterwards I would like to see a report of the results …
blikken terug op 20 jaar DmfA.

Pre-Digital Era

before 1970

- **Information explosion**
  - Growth of **printed** material
- **1939 - Smals founded**
  - Starts with data & information management
- **1945 - "As we may think" essay**
  - Vannevar Bush
  - Memex, a machine to combine lower-level technologies to achieve a higher level of organized knowledge
- **1956 – 1st hard drive**
  - IBM, 5 MB capacity drive

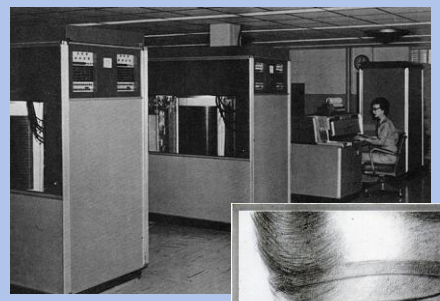

Early Digital Age

1970's – 1980's

- **1970 - Relational Database Model**
  - Edgar J. Codd
- **1981 - IBM PC released**
  - Start of decentralized data generation
- **1984 - Awareness of "information overload"**
  - A. Toffler, "The Third Wave"
  - Managing digital information
- **1986 - Postgres publications**
- **1988 - Postgres prototype**
- **1989 - Postgres v1**

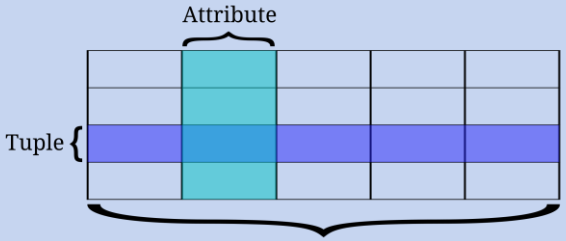

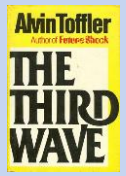

Internet Era

1990's

- **1991 - World Wide Web**
  - Tim Berners-Lee
  - Launches exponential globalized data growth & "***Big Networking***"
- **1994 - Postgres v4.2**
  - Berkeley releases it under MIT License
- **1998 - Google founded**
  - Search engines revolutionize web-scale data retrieval
- **(Re-)Centralization of data**
  - Get rid of data silo's
  - Data Warehousing for analysis

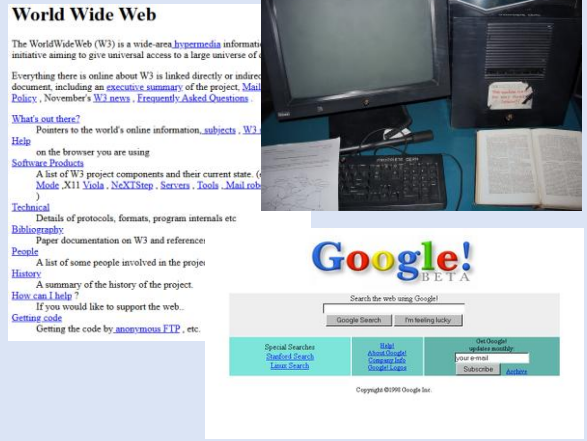

Big Data Beginnings

2000's

- **Explosion of Unstructured Data**
  - Rise of social media
  - Rise of e-commerce platforms
- **2001: Big Data Concept**
  - Gartner defines the "3 Vs of ***Big Data***"
  - Volume, Velocity, Variety
- **2004-…: Technology Advances**
  - MapReduce
  - Hadoop
  - Postgres EDB founded

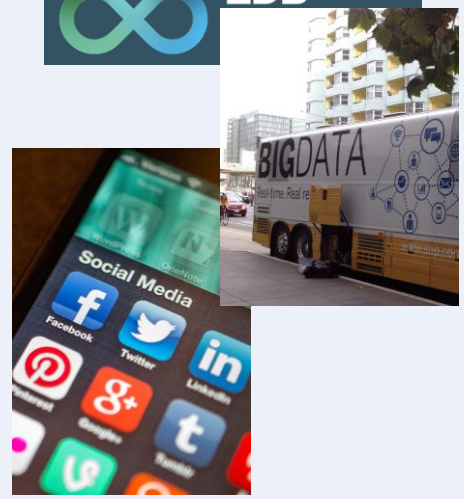

Data-driven Era

2010's

- **Data growth acceleration**
  - "There were 5 exabytes of information created between the dawn of civilization through 2003, but that much information is now created every 2 days." – Eric Schmidt, of Google, 2010
- **AI and machine Learning**
  - AI drives demand for large-scale datasets
- **Real-time processing**
  - Technologies like Apache Spark enable instant insights
- **Cloud adoption takes off**
- **Regulations for data**

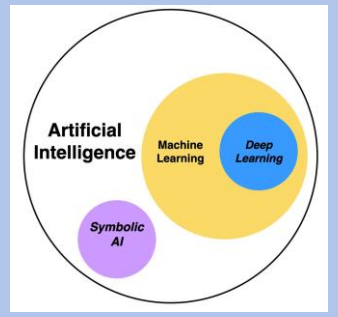

aws

Azure

Google Cloud

Data Overload & Governance

2020's

- **Exponential data growth**
  - Global datasphere predicted to reach 175 zettabytes by 2025 (IDC)
- **Accelerated challenges**
  - Information overload
  - Privacy concerns
  - Ethical AI
- **Governance and regulations**
  - GDPR, Data Act, AI act …
  - Responsible data use
- **Generative AI skyrockets**
- Perfect cocktail of ***Big Compute, Big Networking, Big Data, Big Governance, Big Regulation, …***

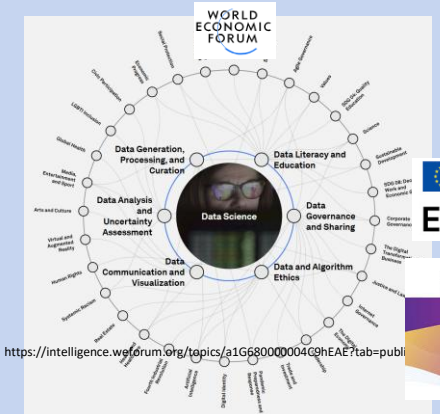

https://intelligence.weforum.org/topics/a1G680000004C9hEAE?tab=publi

EU AI Act

Future?

2030's & beyond

- **AI-Native Data Handling?**
  - Autonomous systems managing and curating data
- **Quantum Data Processing?**
  - Tackling data volume and complexity.
- **Ethical Data Ecosystems?**
  - Increased transparency and control over data usage
- …
- **Reinvent, rethink data & information management?**
  (to catch up with today's challenges)

Compound Annual Growth
Data/Storage
(baseline 2016)

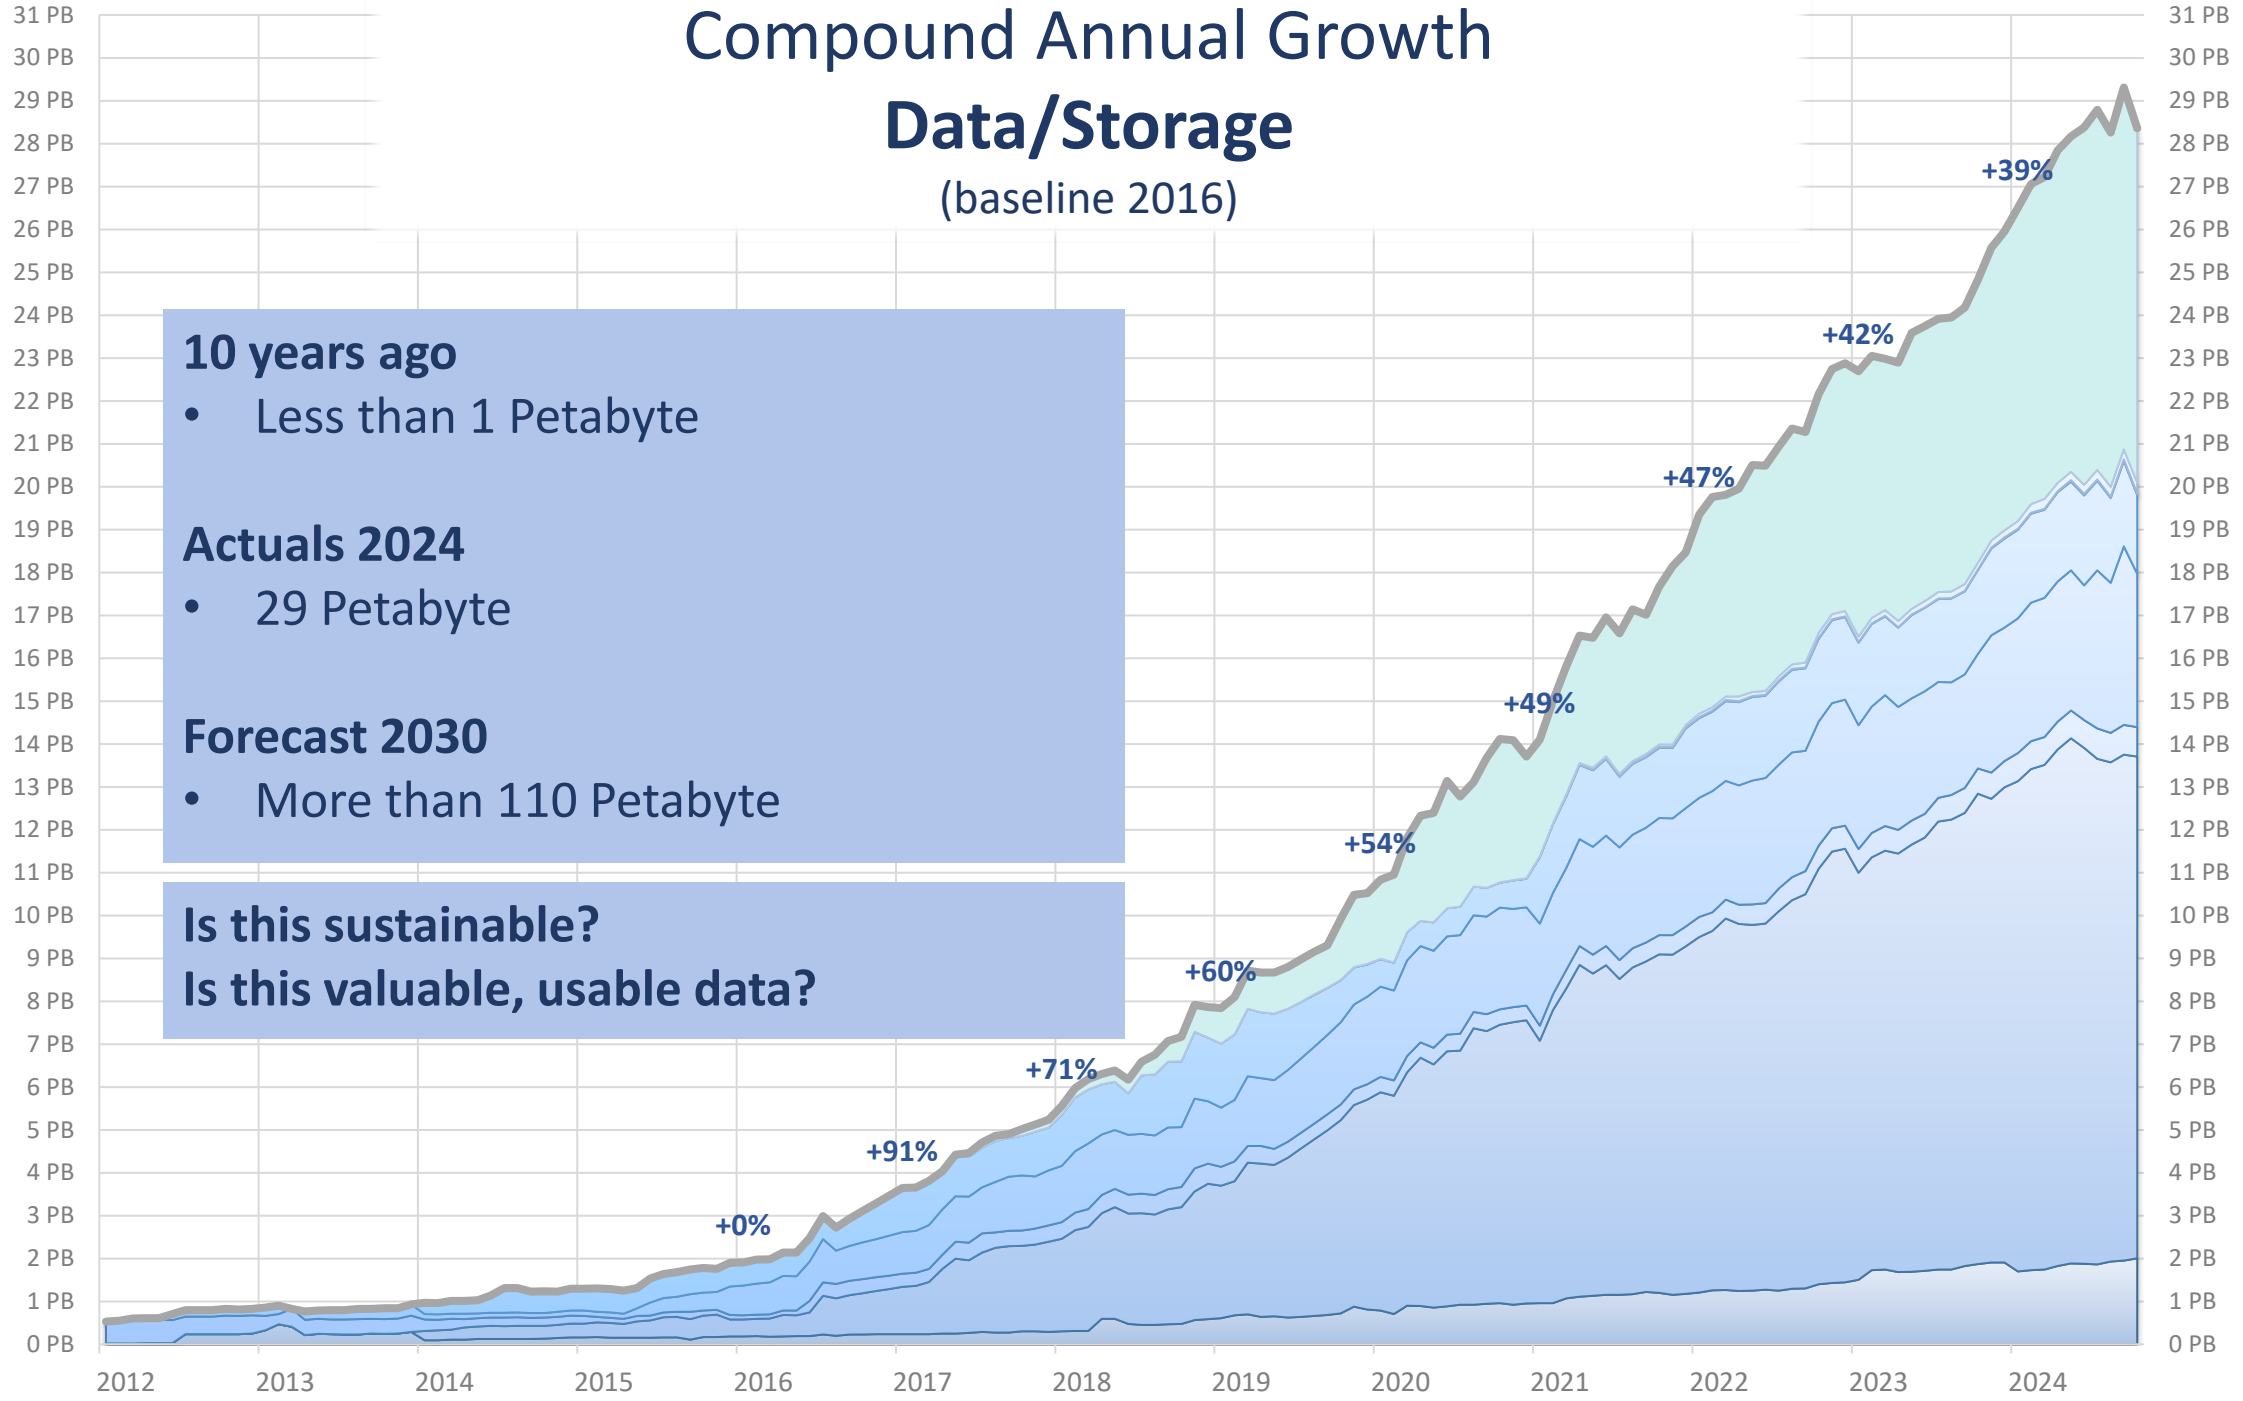

10 years ago
- Less than 1 Petabyte

Actuals 2024
- 29 Petabyte

Forecast 2030
- More than 110 Petabyte

Is this sustainable?
Is this valuable, usable data?

Data & Storage Policy at Smals

1. **Improve awareness**

2. **Keeping global data & storage costs under control**
   - Holistic approach, not just per silo or individual organization
   - Combining a top-down (business → technical) with a bottom-up (technical → business) view

3. **Avoiding 'wasteful' data & storage capacity**
   - Excessive reserved capacity, data/storage duplication, cold & dead data/storage, mismatch quality of service, …

4. **Identify opportunities for rationalization and improved value management**
   - Both on the technical and on the business side
   - Proactive rationalization: prevention is better than remediation
   - Data that doesn't exist costs 0€

Green IT?

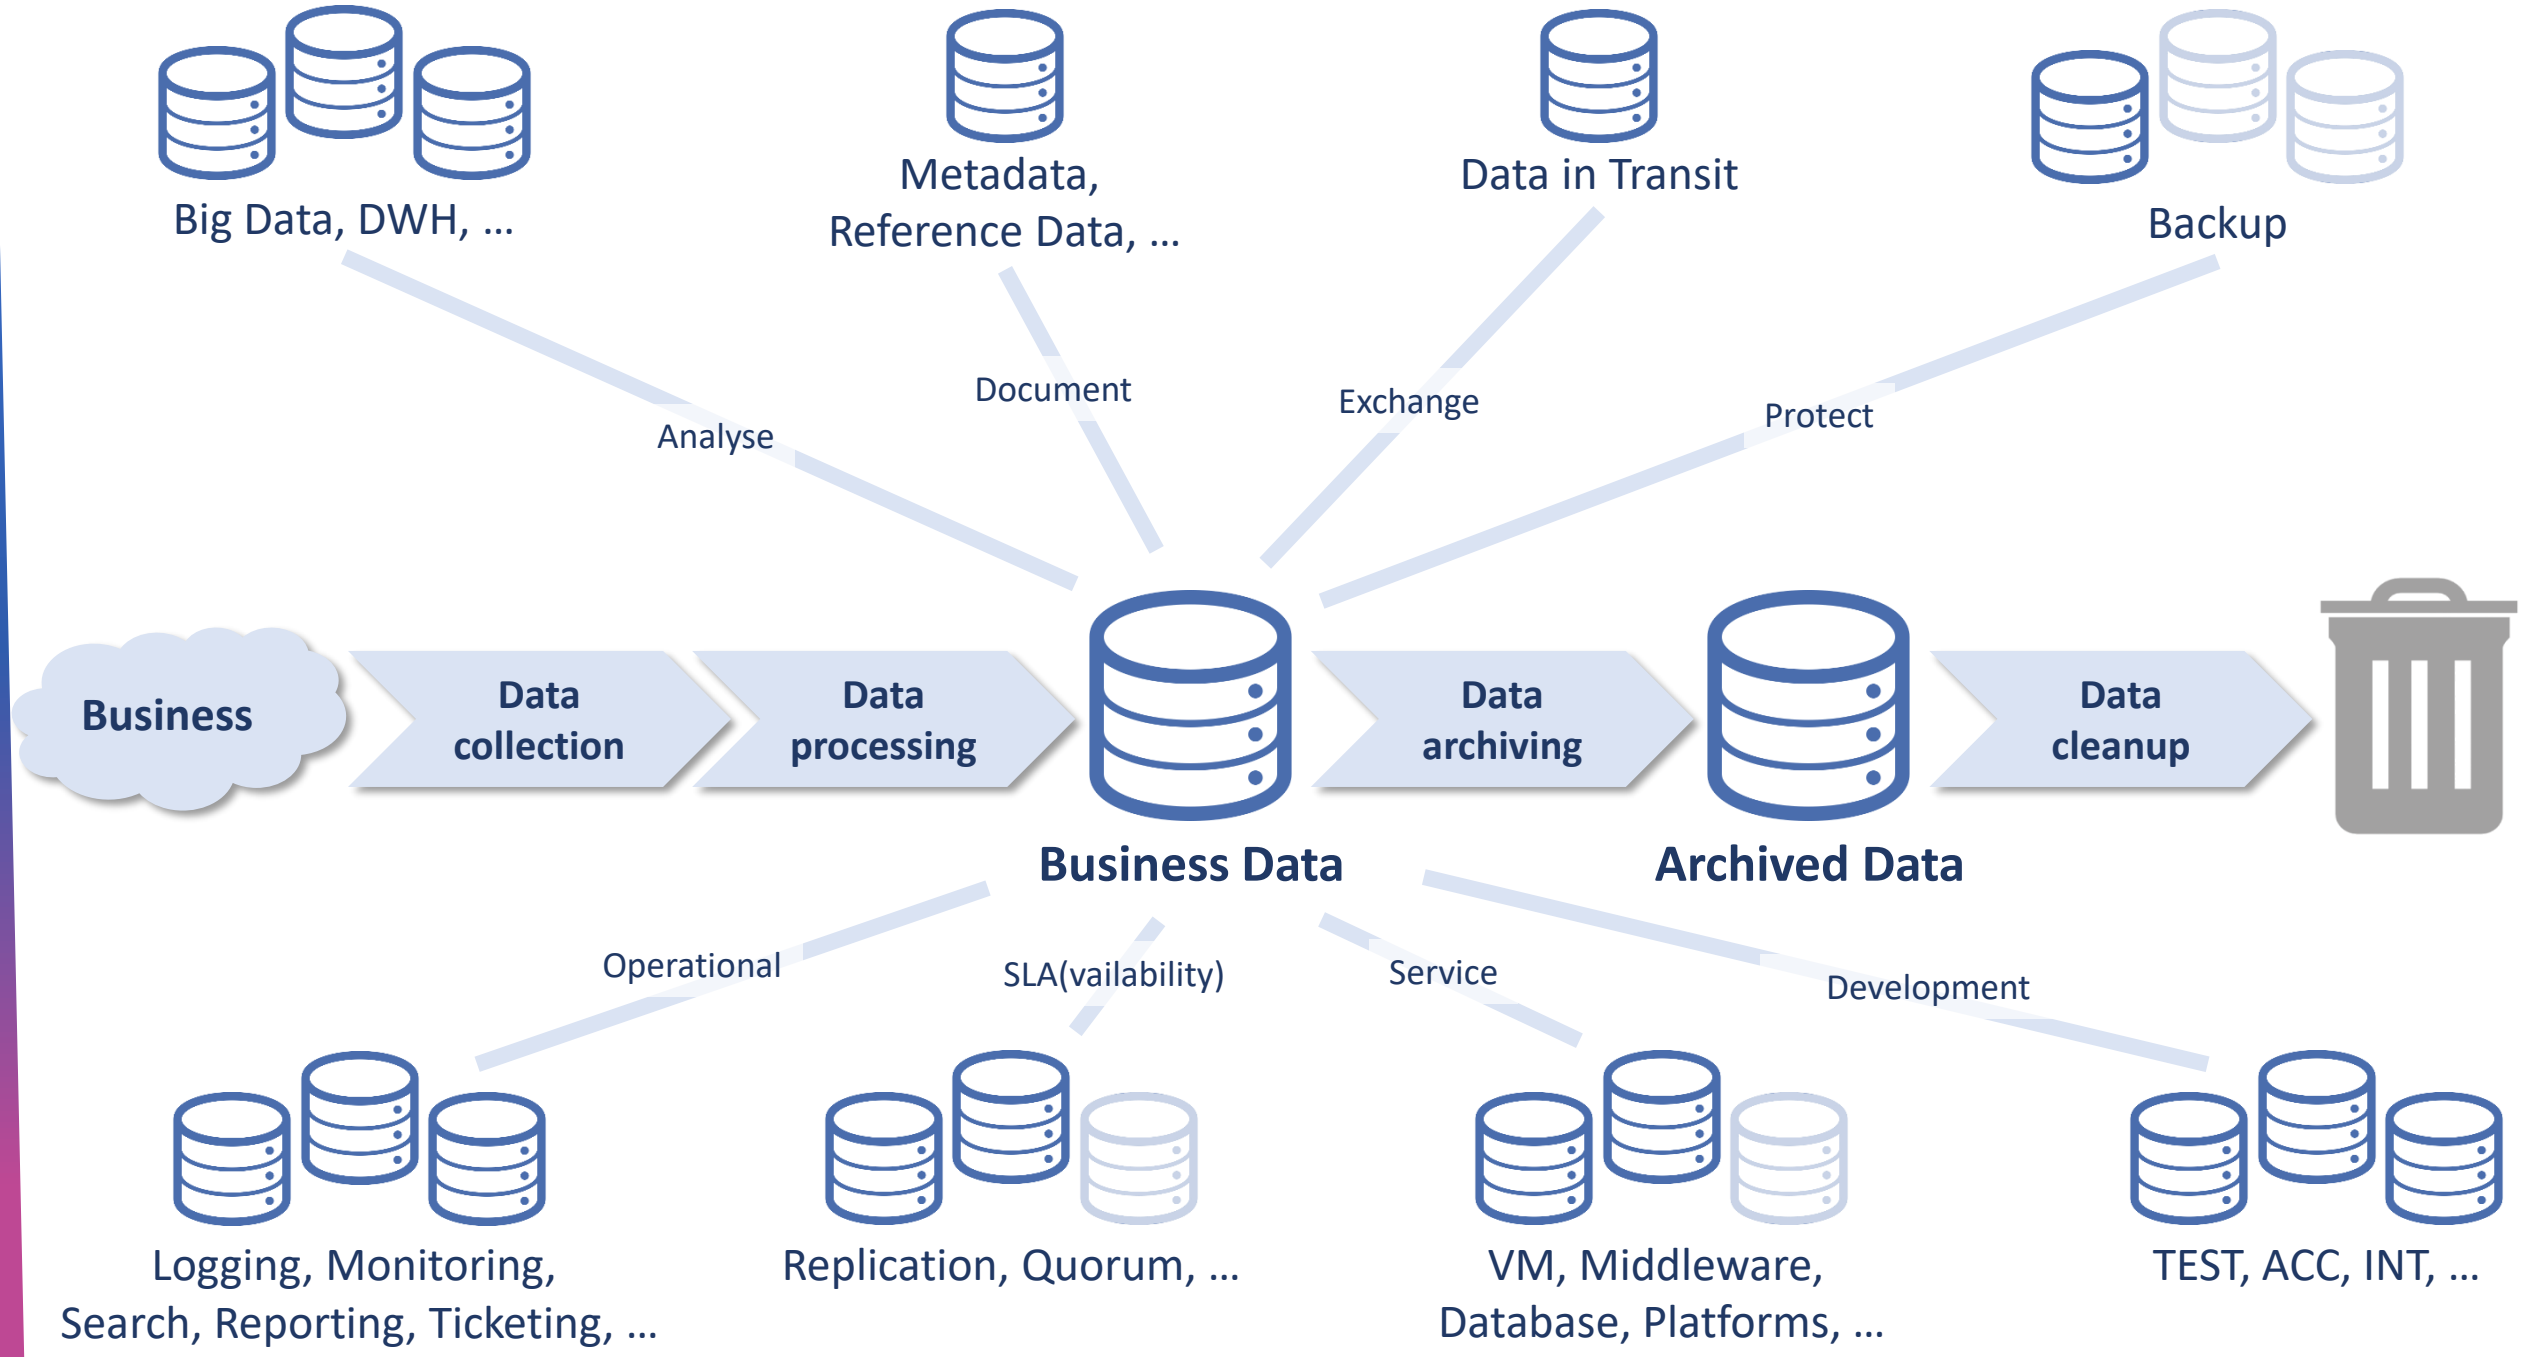
Big Data, DWH, …
Metadata,
Reference Data, …
Data in Transit
Backup
Analyse
Document
Exchange
Protect
Business
Data collection
Data processing
Data archiving
Data cleanup
Business Data
Archived Data
Operational
SLA(vailability)
Service
Development
Logging, Monitoring,
Search, Reporting, Ticketing, …
Replication, Quorum, …
VM, Middleware,
Database, Platforms, …
TEST, ACC, INT, …

1 Vermijden van Data

Bij projecten volop inzetten op API Economy, Referentie Datasets, Authentieke bronnen, Mutualisering, XaaS, …

2 Doorgedreven data analyse

Bij projecten het data landschap volledig in kaart brengen: Retentie, Waarde, Afgeleide Data, Globale Footprint, …

3 Archiveren op goedkopere platformen

Bij projecten de archiveringswaarde en termijn op voorhand bepalen, bestaande databanken transfereren, …

4 Elimineren van Data

Retentie van Business/Technische data zo kort mogelijk houden en elimeren na de houdbaarheidsdatum

5 Business cases Data Analyse & Rapportering

Copieën van data beperken adhv duidelijke ROI opvolging en vooraf bepaalde retentie richtlijnen

6 Data voor Ontwikkeling nauwgezet beheren

Ontwikkelomgevingen (TST/INT/ACC/…) onder controle houden en elimineren of archiveren bij project closure

7 Best Practices voor infrastructuur toepassen

Optimalisatie technieken ten volle benutten en een actief rationaliseringsbeleid voeren

8 Afstemmen SLA met business noden

Beschikbaarheid en performantie fijnstellen adhv een kritische analyse van de reële noden en baten

9 Backups en Archieven van data en systemen

Enkel waar zinvol en waardevol inzetten, duidelijke retentie termijnen vastleggen, elimineren waar mogelijk

10 Optimalisatie op Financieel & Efficiëntie vlak

Technologische standaardisatie, fitness for purpose, proactief rationeel capaciteitsbeheer, …

Data

Search, Reporting, Ticketing, …

Database, Platforms, …

Smals
ICT for society

More than 80 years of data & information management

4.450
Databases

2.200 PROD
2.450 NON-PROD

1.800
Oracle
40,45%

1.543
SQL Server
34,67%

803
Postgres
18,04%

165
InfluxDB

58
ElasticSearch

37
MySQL

25
Cassandra

18
Neo4J

(*) Disclaimer: Count(databases) is eventually consistent ☺

Why is Smals interested in open source?

- It’s free !
- Everyone’s using it!
- It’s better because it’s open!
- It’s more secure because the code is open!
- No vendor lock-in!
- …

Smals
ICT for society

Some Principal values & drivers of Smals

- **Sustainable** ICT solutions and services
- **Community** empowerment
- **Collaborative** innovation
- Overarching “enterprise” architecture **vision** for public sector ICT
- **Quality** driven with a particular focus on **business continuity** and **security**
- Data **ethics** and data protection as a fundamental cornerstone in our approach
- Digital inclusion and **citizen-centricity** as the default
- **Cost-Conscious** and **lean** mindset to maximize **value** for our community …

→ **Synergetic**: giving rise to a whole that is greater than the simple sum of its parts
→ **Enabler**: ensuring indirect positive benefits for society
→ **Invisible:** leaving the spotlight to our members

So why is Smals interested in open source?

- **Alignment with our Vision, Culture, and Mindset**
  - Principal values & drivers resonate with how we look at the ICT world
- **Quality-driven Community**
  - Fostering a quality-driven mindset to ensure operational stability and robustness
- **Open standards**
  - Adoption of open standards to facilitate longevity and to mitigate vendor lock-in
- **Business Continuity Assurance by Embracing ‘Free(dom)’ technology**
  - Utilizing free technology & the freedom to operate
  - Counter ‘Machiavellian Licensing Plots’ and balance the power of proprietary vendors
- **Creative Nexus**
  - It is where innovation magic happens

So why is Smals interested in open source?

- We follow at the same time a **strategic** and an **opportunistic** approach
- We continuously balance proprietary vs commercial/hybrid vs community-driven
- We try to find the sweet spot between different objectives
  - Financial & procurement considerations (cost, scale of deployment, tendering, …)
  - Enterprise risk mitigation (vendor lockin, business continuity, …)
  - No-brainer (linux, core tools, drivers, …) + Technological motivations
  - Innovation & short time to market …
- When we follow a dual vendor approach, open source is typically the alternative
  - Lack of enterprise “features” sometimes hampers adoption (as well as TCO for commercial/hybrid!)
- We ~~think~~ ~~believe~~ hope open standards will bring us 99% of portability (on premise)
  - Or at least more than 0%
  - Proven in the past that proprietary → open source is a lot harder than the other way round

We definitely want to see an increase in its usage,
employing an embedded ,disciplined, and methodological approach

Asking “why open source” is a relevant, but strange question if you consider…

Open Source is everywhere @World

- It depends on what you include:
  - Products, Platforms, Tools, Libraries,
    Drivers, Code snippets, …
- **OSSRA report finding (2024)**
  - Open source components and libraries
    are the backbone of nearly every application across all industries
  - **96% of total codebases contain open source**

96%

of the total codebases
contained open source

Proprietary vendors would not exist without open source !

Smals
ICT for society

15 minute braindump…

Hence Open Source is also everywhere @Smals

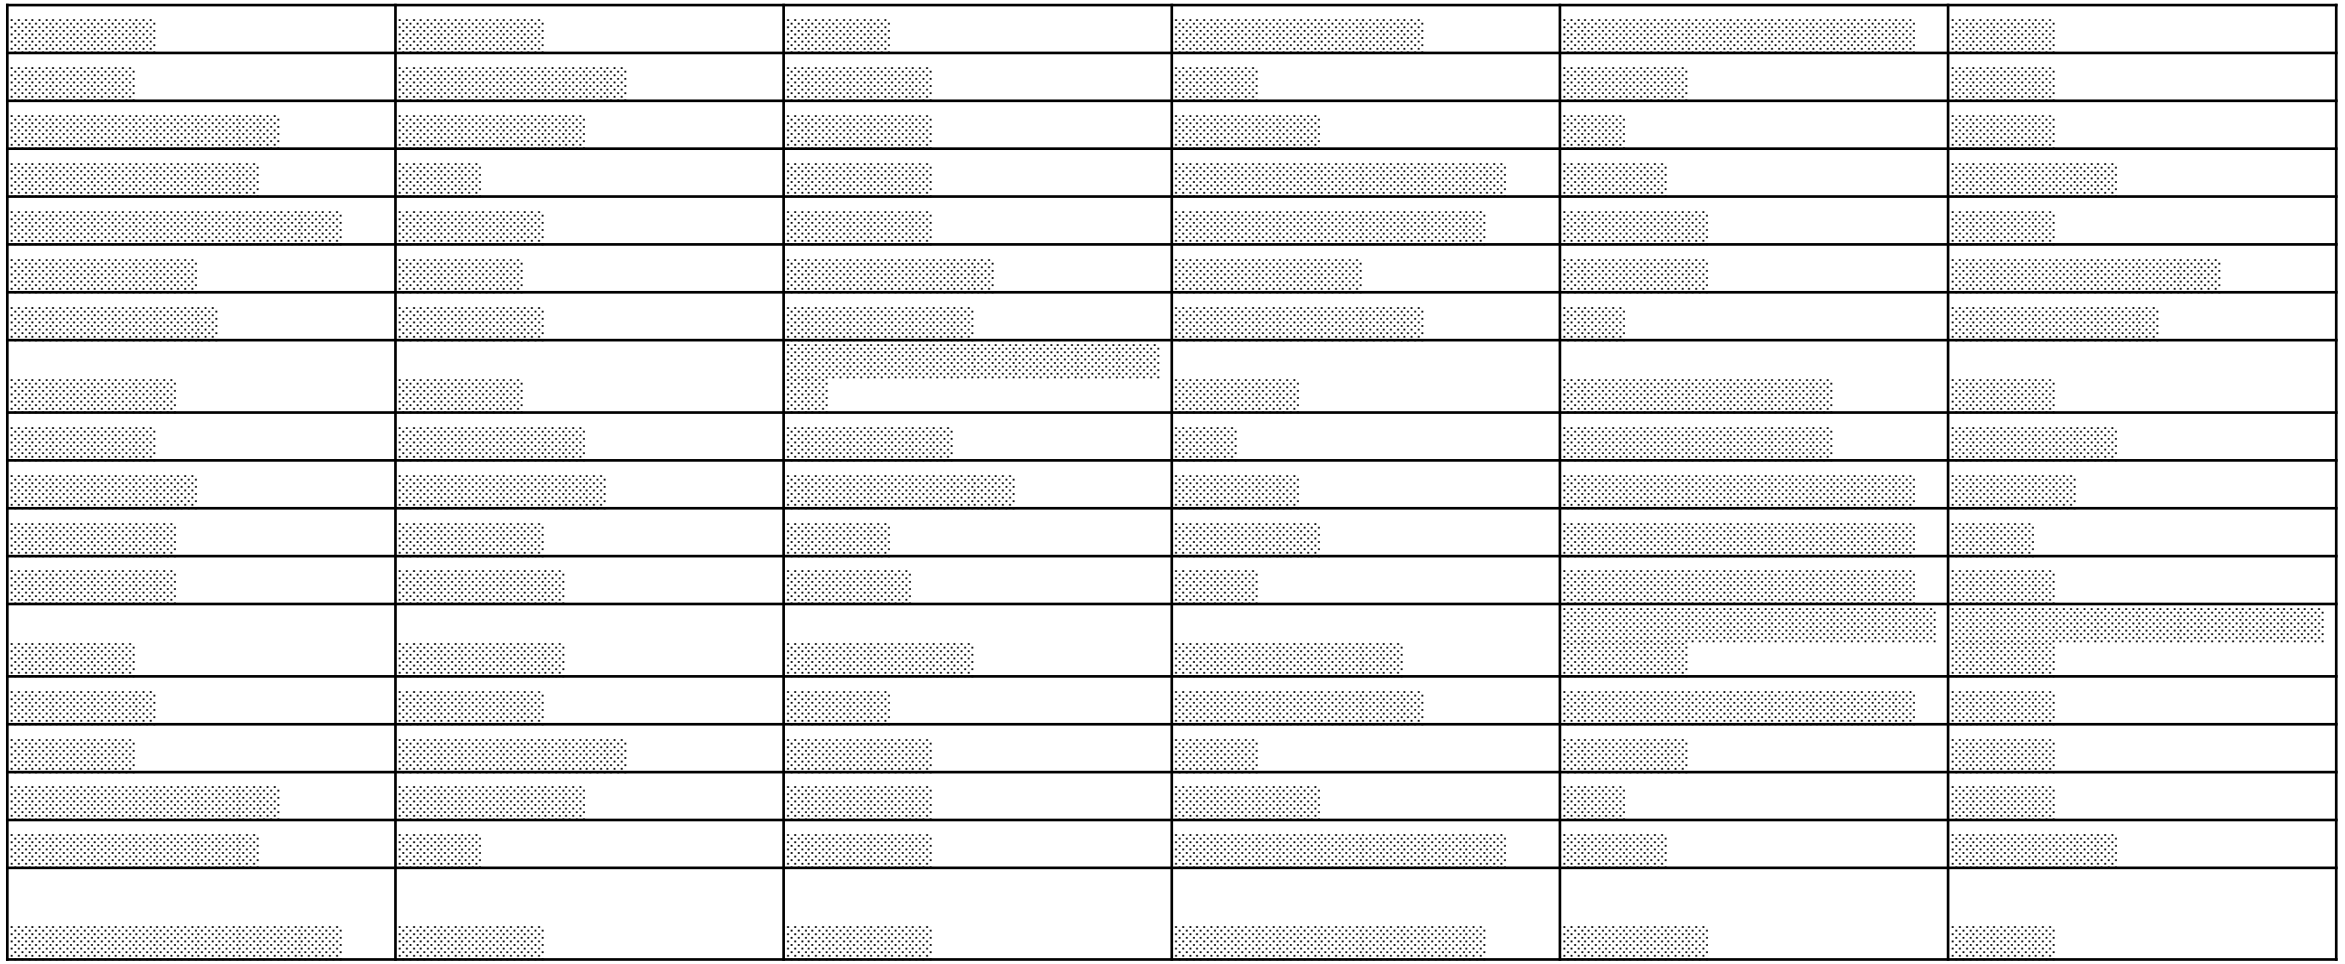

…

Apparently, we <u>don't always require</u> enterprise versions/support

Open Source is everywhere at Smals

But for some platforms we do <u>want</u> enterprise versions/support

Red Hat

Linux, Openshift, JBoss, AMQ,…

EDB™

POWER TO POSTGRES

Smals

ICT for society

Postgres products used by Smals

- EnterpriseDB Advanced Server (EDB-AS)
- Postgres Community version (small footprint)
- Postgres Enterprise Manager 9.5 (PEM)
- EDB Failover Manager 4.4 (EFM)
- Backup and Recovery Tool 2.6 (BART) (will switch to Commvault)
- PG_DUMP / PG_RESTORE
- Foreign Data Wrapper – Postgres / Oracle (FDW)
- EDB Migration Toolkit (EDBmtk)
- LiveCompare

803

Postgres

18,04%

Smals
ICT for society

Postgres Services delivered by Smals

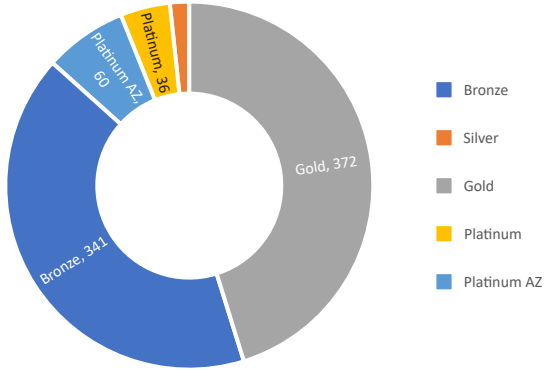

Databases Management

Service description	Relational and NoSQL databases provision & management					
Databases Services Catalog	**Database as a Service (DBaaS)** Self Services		**Fully Managed Database** DBA Services			
Self Services / Standard Services / Delivery Services / Value-Added Services	**Self**-Managed	**Cloud**-Managed	Managed **Basic**	Managed **Standard**	Managed **Advanced**	
Release number : 4.6 Release date : MAY-2023	Self-provisioning | Self-admin	Self-prov | Admin facilities | DBA assistance	Operational support delivered by DBA	Operational support delivered by DBA	Operational support delivered by DBA	Operational support delivered by DBA
	For G-Cloud Customers	For Smals	For Smals	For Smals	For Smals	For Smals
Supported DB-Environments						
Dev | Test | Int	✓	✓	✓	✗ Not available	✗ Not available	✗ Not available
Acceptation	✓	✗ Not available	✓	✗ Not available	✗ Not available	✗ Not available
Simulation	✓	✗ Not available	✓	✗ Not available	✗ Not available	✗ Not available
Production	✓	✗ Not available	✓	✓	✓	✓
Supported DB-Engines						
Database Matrix / RDBMS / EDB POSTGRES	✗ Not available	✓	✓	✓	✓	✓
RDBMS / Microsoft SQL Server	✗ Not available	✗ Not available	✓	✓	✗ Not available	✗ Not available
RDBMS / MySQL	✗ Not available	✗ Not available	! Deprecated	! Deprecated	✗ Not available	✗ Not available
RDBMS / ORACLE DATABASE	✗ Not available	✗ Not available	✓	✗ Not available	✓	✓
NoSQL / elasticsearch	✓	✓	✓	✓	✗ Not available	✗ Not available
NoSQL / influxdb	✓	✗ Not available	Q4 - 2023	✗ Not available	✗ Not available	✗ Not available
NoSQL / Neo4j	✓	✗ Not available	Q4 - 2023	✗ Not available	✗ Not available	✗ Not available

Setup for Postgres EDB Silver SLA @ Smals

No standby configuration
If master fails
    a recovery or restore will be needed
Intervention manual

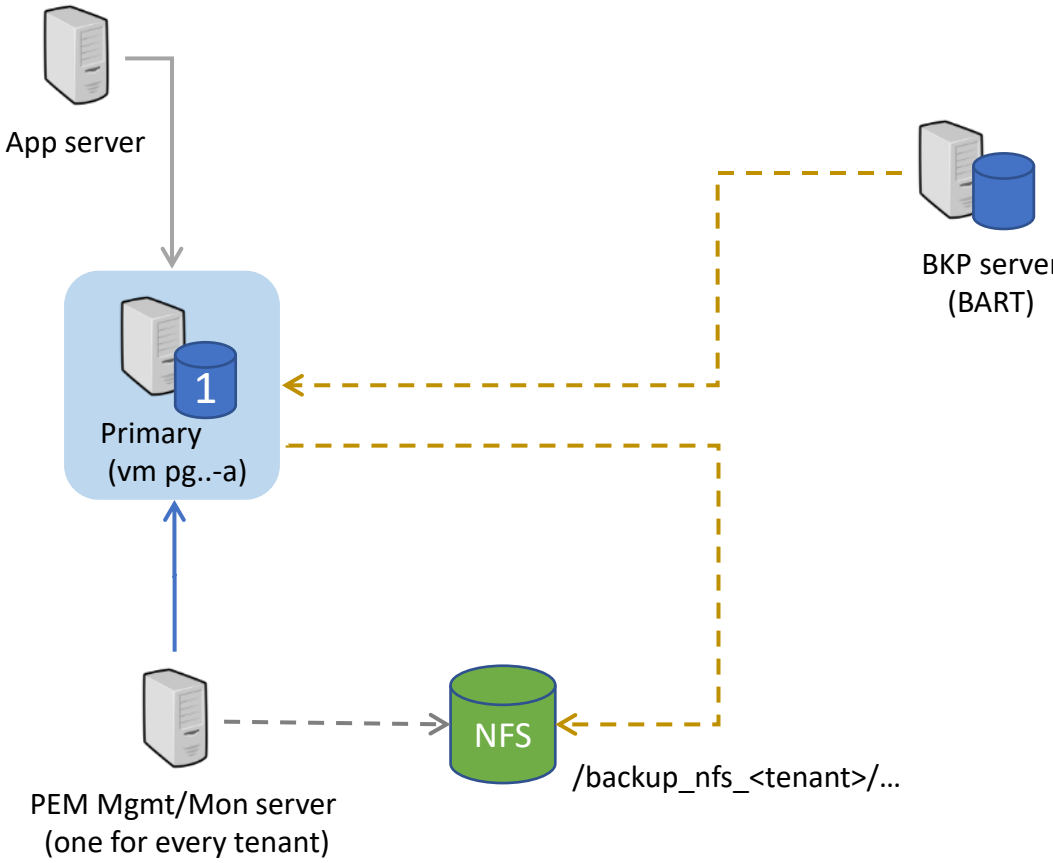

Data center 1

Data center 2

Data center 3

Communication when the Master is running

X

VM with postgres
High perf storage

VIP

VIP managed by DBA

Setup for Postgres EDB Silver SLA @ Smals

One VIP per Postgres Cluster (port)
EFM will NOT failover to standby
Failover: MANUAL

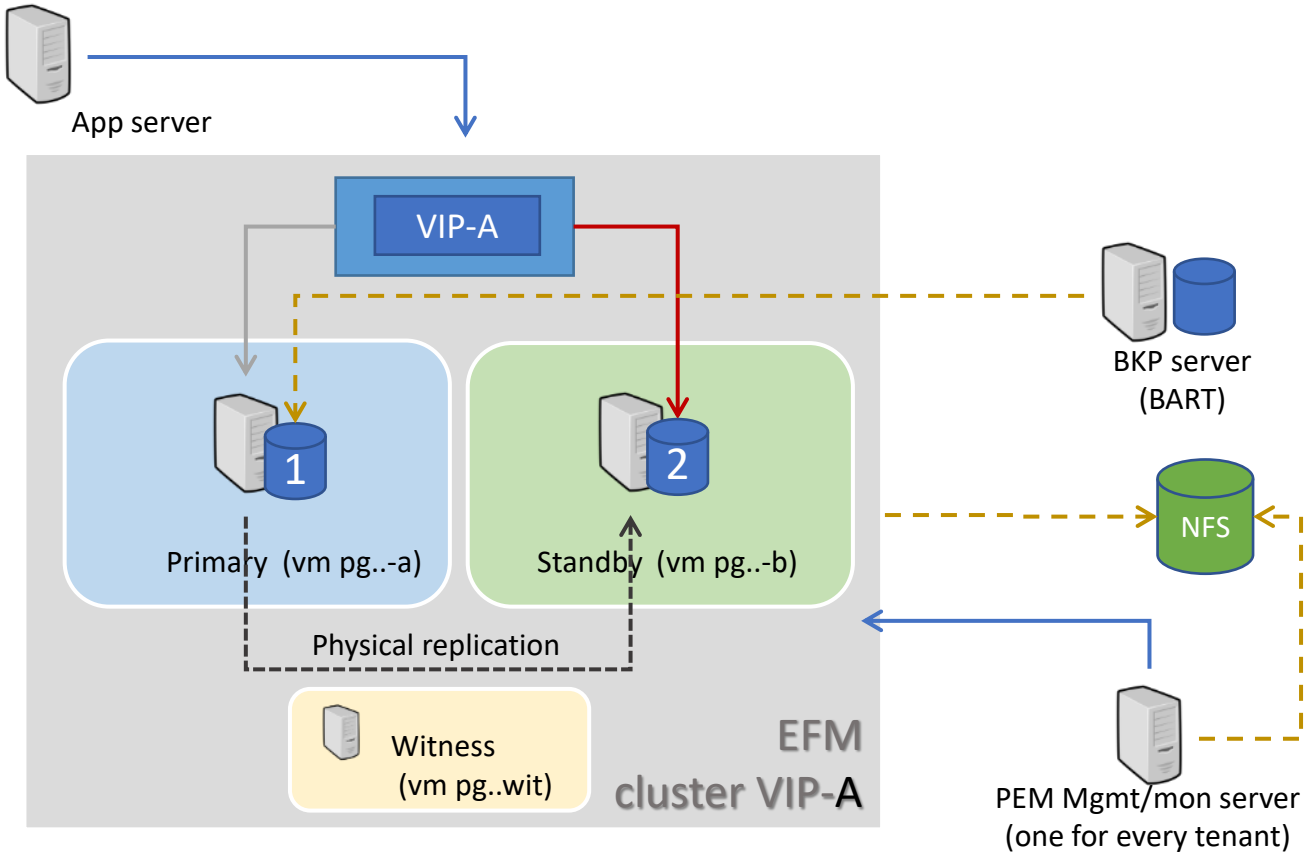

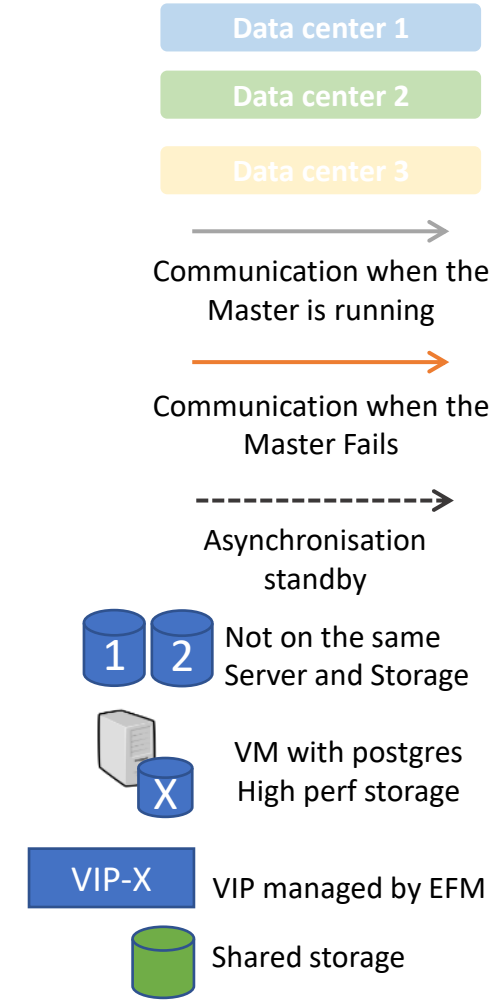

Setup for Postgres EDB Gold SLA @ Smals

One VIP per Postgres Cluster (port)
EFM will failover to standby
Failover: automatic – 5min

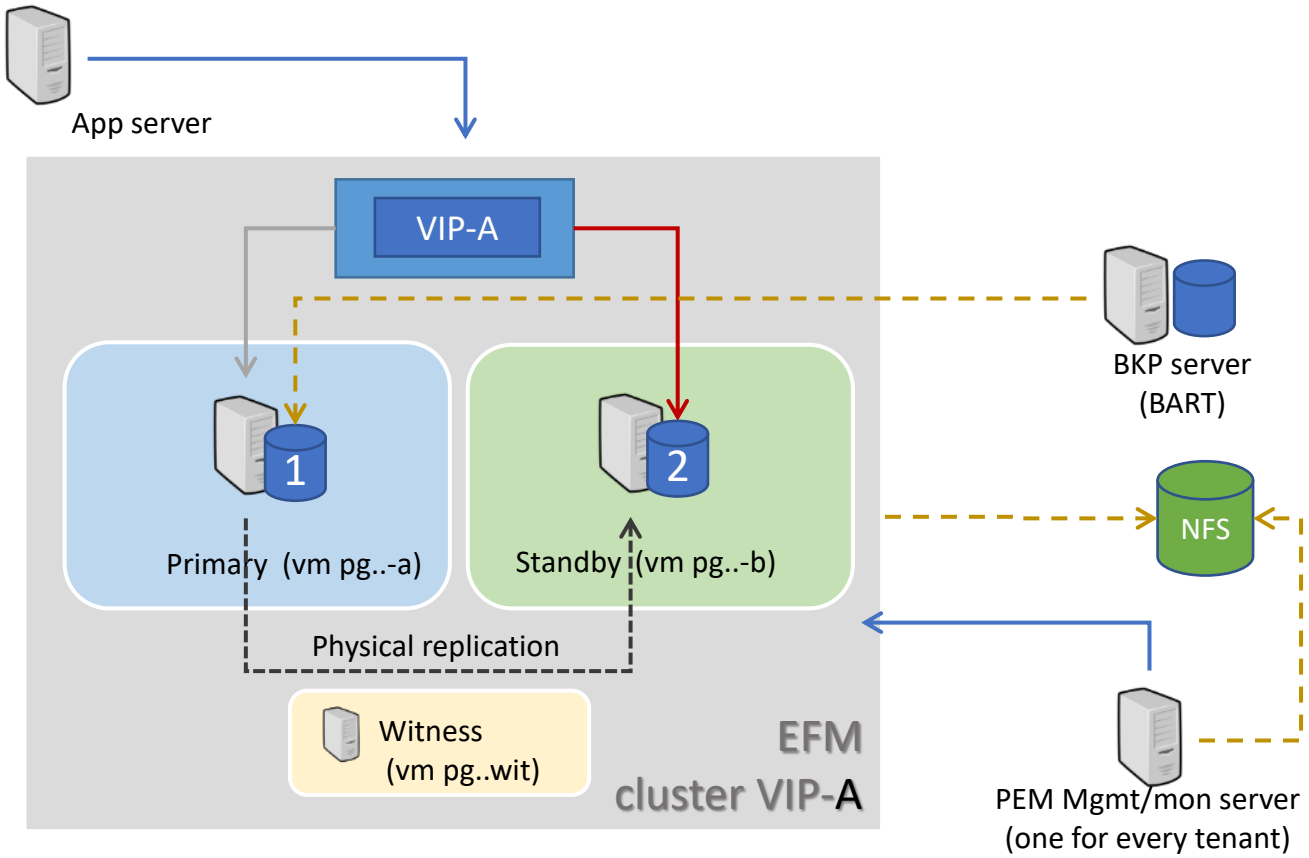

Data center 1

Data center 2

Data center 3

Communication when the Master is running

Communication when the Master Fails

Asynchronisation standby

1 2 Not on the same Server and Storage

X VM with postgres High perf storage

VIP-X VIP managed by EFM

Shared storage

Setup for Postgres EDB Platinum SLA @ Smals

One VIP per Postgres Cluster (port)
EFM will failover to standby
Failover: automatic – 1min
One standby is in SYNC with primary

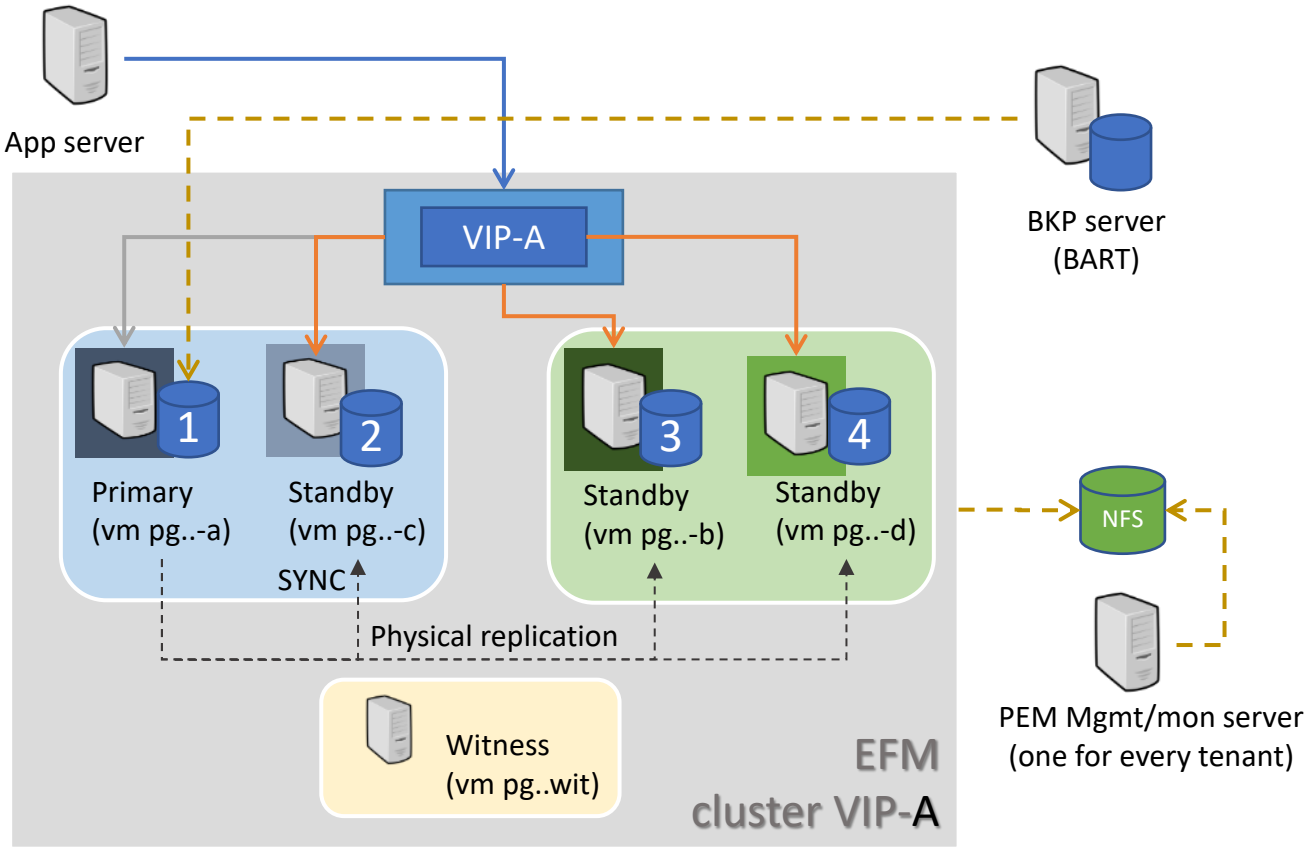

Data center 1

Data center 2

Data center 3

Communication when the Master is running

Communication when the Master Fails

1 2

Not on the same storage
Class 1 disks

VM on different hardware

VIP-X

VIP managed by EFM

Shared storage

Setup for Postgres EDB Platinum Multi-Availability Zone SLA @ Smals

Two VIP's per Postgres Cluster (port)
Failover: - within the same availability zone: automatic – 1min
- availability zone: manual via AZMA
One standby is in SYNC with primary

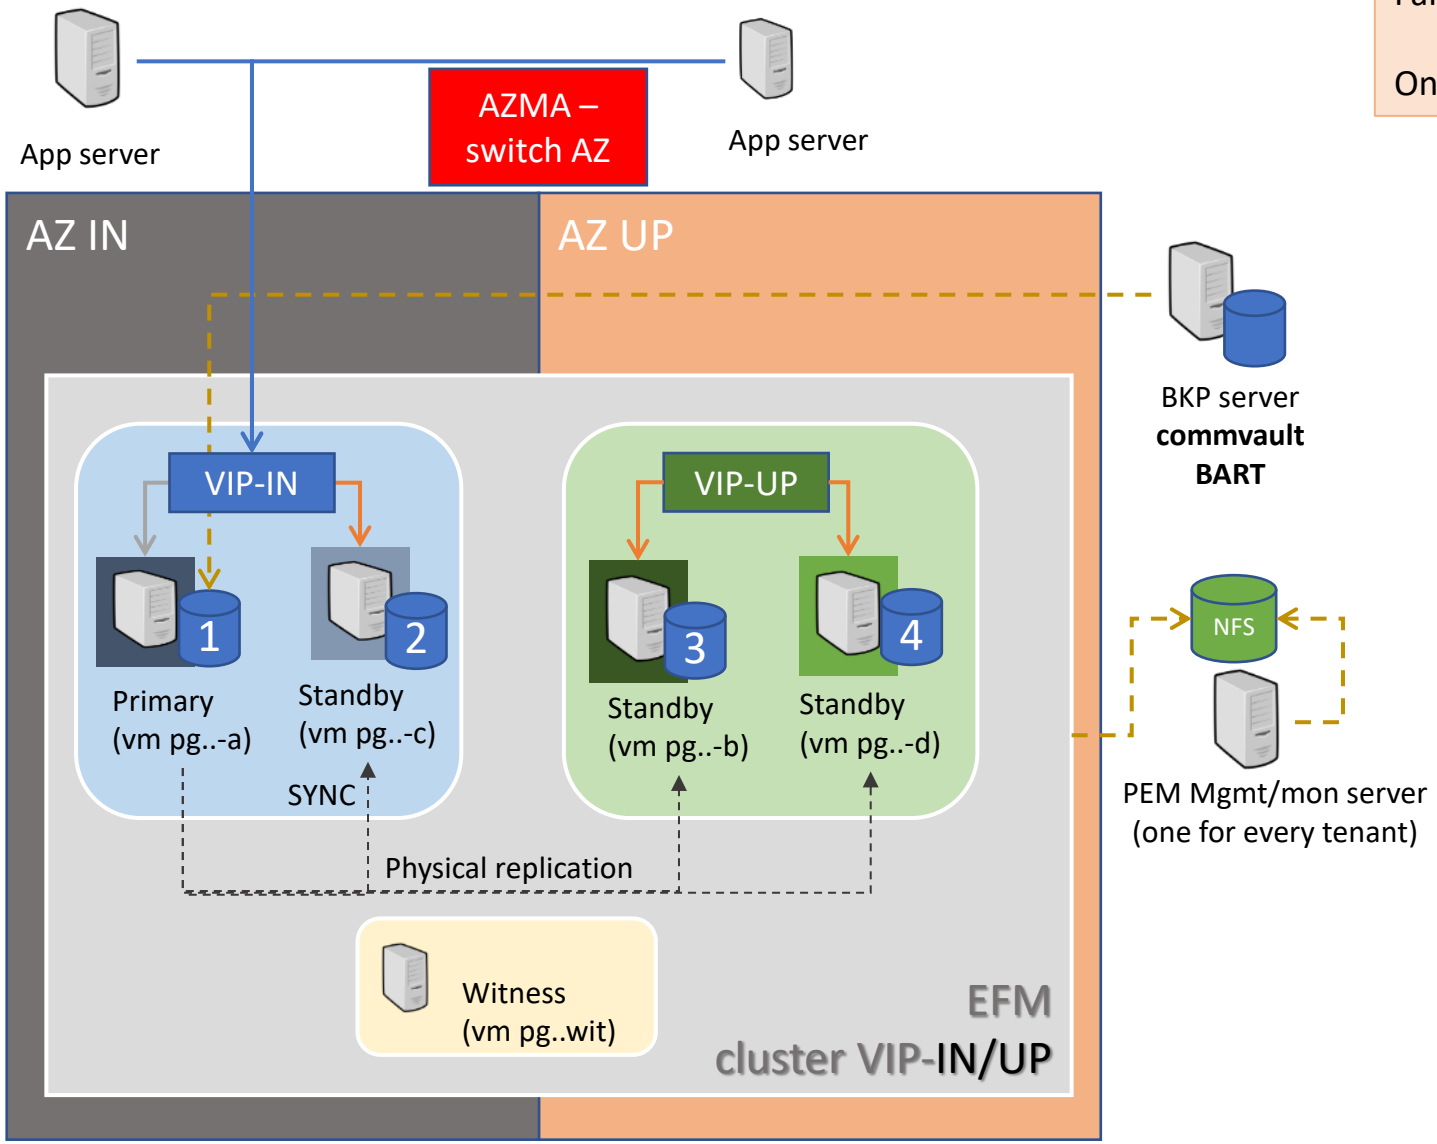

Data center 1

Data center 2

Data center 3

Communication when the Master is running

Communication when the Master Fails

1 2 Not on the same storage Class 1 disks

VM on different hardware

VIP-X
VIP managed by EFM

Shared storage

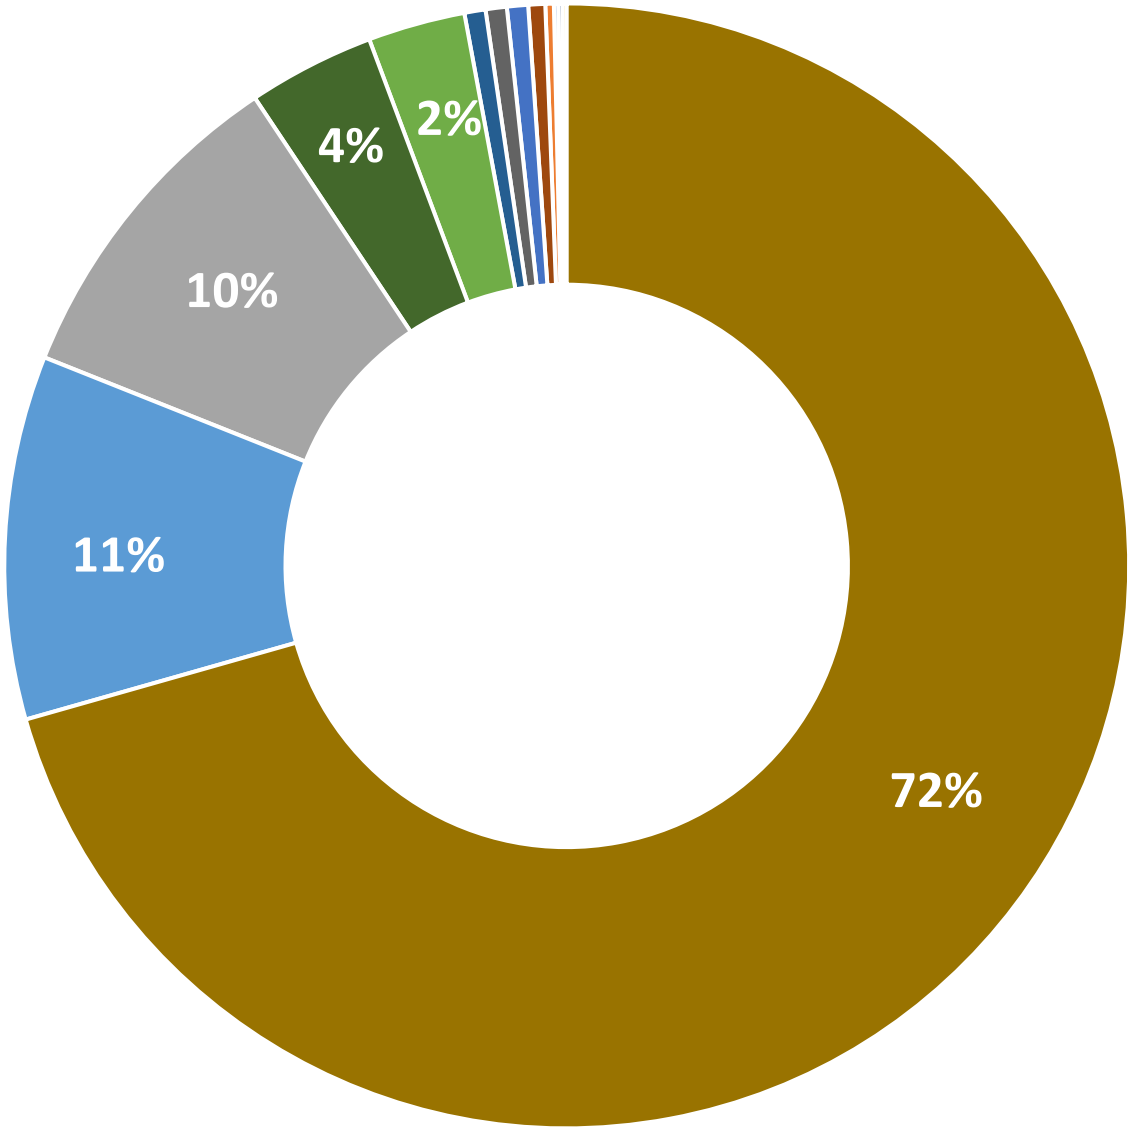

PostgreSQL Versioning

Open-source agility with frequent releases to address community-driven innovation
→ Yearly majors, quarterly minors

- Lifecycle Management target
  - Max 3 versions supported internally
  - 80 / (15 + 5) rule as an objective
- Actual status for latest in-house majors : 76% / (11% + 2%)
- Rhythm is quite challenging, target 1 out of 2 major versions
- Migrations are always ongoing
- Automated upgrades as much as possible (non trivial)
  - Minimize impact: no downtime, no changes, …
  - Target: even during business hours

CloudNativePG - Postgres Operator for Kubernetes

- Planned rollout in 2025
- Will bring many benefits (availability, upgrade, …)

PostgreSQL Lifecycle Management

1. **Logical Replication**
   - Streams changes to the new version.
   - Pros: Minimal downtime, selective data replication
   - Cons: Schema compatibility required; no full DDL support
2. **pg_dump/pg_restore**
   - Logical backup and restore process.
   - Pros: Simple, cross-version compatibility.
   - Cons: High downtime, slow for large datasets, needs extra disk space
3. **pg_upgrade**
   - In-place upgrade of data files.
   - Pros: Fast, handles large datasets well.
   - Cons: Requires app downtime, both binaries needed
4. **Replication Tools (e.g. HVR)**
   - Enterprise-grade data replication
   - Pros: Robust, advanced features
   - Cons: Licensing cost, steeper learning curve, and doesn't always do what's on the sales slides
5. **Snapshot/Restore**
   - Snapshot old DB, restore and upgrade
   - Pros: Simple for small systems
   - Cons: High downtime, not scalable

- **Multiple techniques** are employed
- **Difficult to find the sweet spot** between
  - complexity, downtime, rollback time, DB size, proxy @ application level,…
- **Difficult to automate a "one-way-fits-all" approach**

Postgres vs Oracle?

- **Simpler setup, less configuration, faster to get it up & running**
- Optimizing Oracle licensing requires a PhD
- **Open Source**, so in theory not limited to 1 vendor
- **Postgres tooling is on par** with traditional enterprise toolkits
  - “Marginal” differences, but works as intended
  - OEM = PEM, Rman = Bart, Dataguard = Slave, …
  - Community drives **innovation and tooling**
- **DML-DDL transaction management is more flexible**
  - You can include more in 1 single transaction
  - Auto-commit on DDL operations
- **Backup/restore** has some limitations (*)
  - Instance with multiple DB’s: single DB restore not possible
  - Slow Point in Time Recovery for large DB/high transactions (e.g. 5h30 for 1TB of Write Ahead Logs)
  - (*) PG17 improved pg_basebackup to incremental
- **Bloat: updating a record creates a new version of it**
  - Risk for disk space explosion and performance degradation
  - Needs to be monitored, do a full vacuum
    - Performance impact : Table lock, high I/O, index rebuild

Postgres vs Oracle : Switch now?

- **Strategic directive** in place for **new applications**: PG as default
- EDB supports **compatibility mode** to facilitate migrations
  - Balance between open source and “vendor lockin”
  - Do you really need it / use it?
- **Migration** initiatives remain **difficult**
  - Oracle on baremetal hampers fast ROI for migrations
    - You cannot decommission before all is gone
  - Data architecture & practices are not always today’s way
  - Goes beyond infra/ops work
  - Requires deep implication of development teams
  - Release planning & time & budget

- **Technical effort and complexity**
  - **Oracle PL/SQL stored procedures & triggers** (90% compatible)
  - **Oracle proprietary data types** (Number, Date, Timestamp, …)
  - Subtle differences in **handling of storage structures** (tablespaces, clusters, materialized views…)
  - **Query syntax differences** (e.g. connect by vs with recursive)
  - **Index and performance tuning** …

Overall we remain convinced that Postgres was an excellent strategic choice !

Ok, so Smals is a top Postgres contributor?

- **Our core business is not “technology product development”**
  - We (ab)use technology to (retro)fit the needs of our context
  - Like most organisations we are mainly on the consumer side
  - Abstractions that are (re)usable for the rest of the world pop up occasionally but are not the norm
- **Practical open source contribution in general is not for everyone**
  - Less than 1% of the population is capable to directly contribute in a meaningful way
  - Coding competency, paying with your time, figuring your way around the code base and culture, …
- **We do pay others to contribute for us**
  - e.g. RedHat, Postgres EDB, …
- **Contributions in any form remain a major challenge**
  - In general, not limited to open source
  - Even for in-house contributions to our own transversal micro-cosmos of libraries, platforms, …
  - It’s no free(dom) lunch

We do however contribute to a vibrant local micro-community…

Re Use

BA 4 GOV

g-cloud

Smals ICT for society

What’s on our roadmap?

- **Latest PG version**
- **Bart - Barman - Commvault**
- **CloudNativePG (Kubernetes-native Operator)**
- Transparant Data Encryption (TDE) - SSL
- PG_BOUNCER (Connection Pooling MW)
- Implement fast & robust upgrade method (pg_upgrade, logical replication)
- Follow-up Postgres versions (PG18 – PG20…)
- Very large databases & Postgres
- Monitoring & Tuning
- Multi-master solutions

Smals ICT for society

Key takeaways…

- **We are in the midst of a perfect IT storm that shakes things up**
  - Big networking, Big data, Big compute, Big governance, Big regulation, … Big budget & Big time? ☺
- **Open source & data are the invisible engine of our digital society**
  - We must focus on making it visible (again)
  - Not only to promote it, but also to know what's inside this trusted engine
  - We must drive digital eGov transformation by adopting a strategic data vision and future proof data architecture
  - We must rethink and reinvent the way we work to face the challenges ahead/today
- **Micro-communities can make a difference**
  - We must tap into the collective intelligence of the public sector community to deliver better services to citizens
  - It's about building bridges; we are all part of the public sector family: **Together we stand strong(er)**
- **Postgres is a valid alternative for heavy duty, highly available, high throughput, robust & performant systems**
  - But like all IT you still need to know what you are doing
  - For new developments it is an excellent choice
  - For migration of older / existing systems a careful examination is required

Transform data & information management
to survive today and prepare yourself for the new era

Thank you !

If you have any questions, do not hesitate to contact me

dirk.deridder@smals.be

gcloud.belgium.be

ict-reuse.be

www.ba4gov.be

www.smals.be

Smals
ICT for society