

NATIONAL DATA STORAGE DATA-LAKE-LIKE INTEGRATED DATA MANAGEMENT SERVICES FOR R&E



Data mgmt team:
Krzysztof Wadówka,
Adam Prycki,
Eugeniusz Pokora,
Filip Blicharczyk,
Jan Bróździak,
Krzysztof Błoniarz,
Maciej Brzeźniak

Collaboration (security team):
Mikołaj Dobski (security)
Gerard Frankowski (security)

Collaboraiton (cloud, net, DC):
Jacek Kochan (cloud)
Sebastian Petruczynik (net)
Rafał Mikołajczak (DC)

Director of Data Processing
Technologies Division:
Dr Norbert Meyer

WHO WE ARE





ABOUT PSNC

- **PSNC - mission and activity:**
 - HPC CENTER
 - NREN - PIONIER - CONNECTED TO GEANT
 - R&D center - infrastructure & services projects
 - Services provider:
 - computing, storage, network, hosting, security
 - Certified: ISO 9001/27001
- **PSNC - in numbers:**
 - **400+ people** in 4 divisions: computing, applications, network services, network
 - **HPC: 6 PFLOPs -> 36 PFLOPs** (2023)
 - **Cloud: 12 000** cores, **80 TB** RAM + 200 servers (2023)
 - **Storage: 80 PB** (disk) + **20 PB** (tape) -> **0.5 ExaB** (disk + tape)
 - **Network:** 11 000 km of fibers in 21 cities - WAN & MAN network



PSNC - POLISH NREN AND SERVICES PROVIDER FOR HPC/CLOUD/STORAGE APPLICATIONS

Polish Optical Internet Research Centre

Headquarter

Labs and offices

Primary
Data Centre



Living Labs

Co-working area, spin-off



Metropolitan Area Network POZMAN


100 km of fibre optic cables
10 institutions connected



Secondary Data Centre



ABOUT PSNC



Poznan University of
Technology Campus

My office ;)


Centrum Badańcze
Polskiego Internetu...

Data Center

SPEAKERS - PSNC' DATA PROCESSING TECHNOLOGIES DIVISION

- **INTRO, National Data Storage PROJECT**

- **Maciej Brzeźniak**
- **maciekb@man.poznan.pl**



- **SECURITY USE-CASE (DATA-LAKE)**

- **Mikołaj Dobski**
- **mikolajd@man.poznan.pl**



*DISCLAIMER: we're only fronting the work of many people
within our team, in PSNC and in partnering institutions*

CHALLENGES

IN (SCIENTIFIC) DATA MANAGEMENT




CHALLENGES IN DATA MGMT: VOLUME

2015 PREDICTIONS:

- growing **volume**: PetaBytes
- pressure for **performance**: GB/s, IOPS
- user expectation: **ease of use**

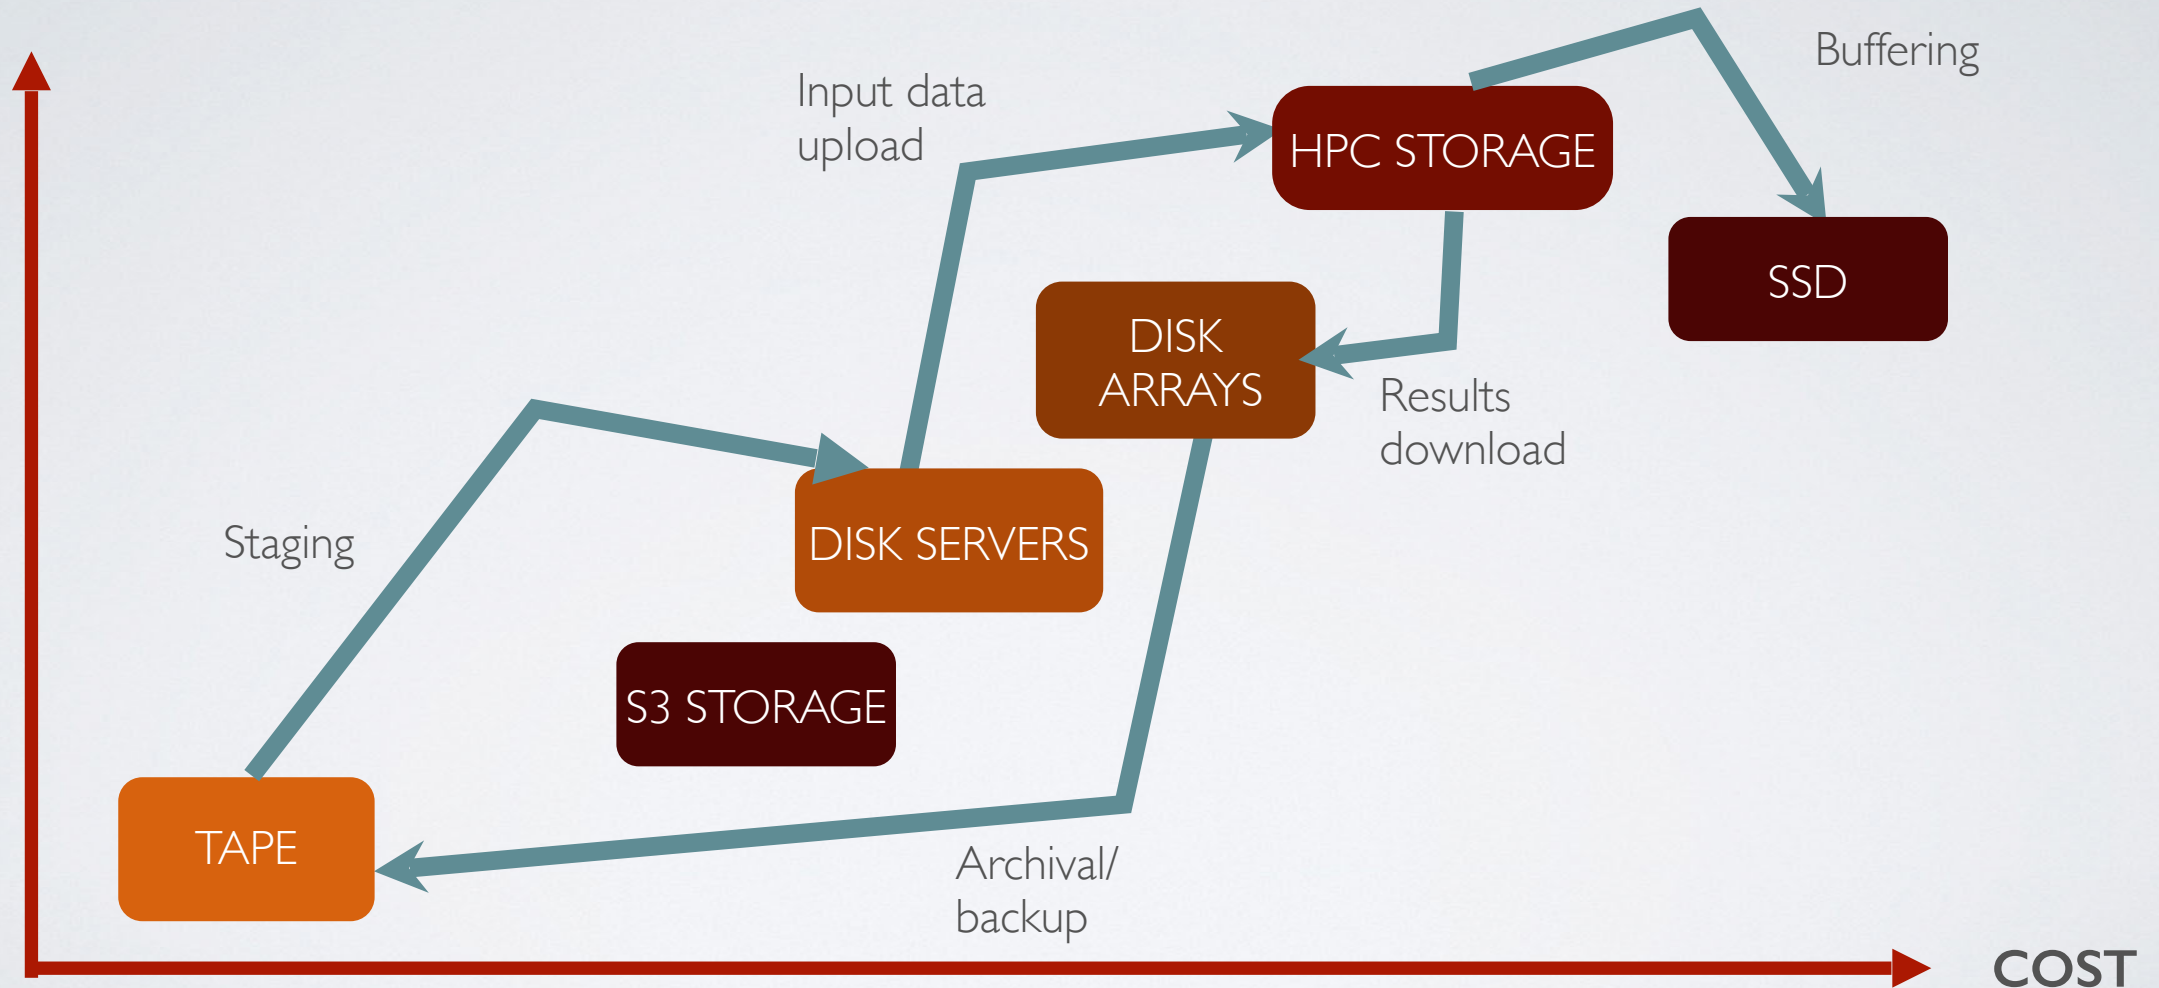
SEEN IN 2023:

- Internet users growth is faster than global population growth!
- Most of data produced in the Internet



CHALLENGES IN DATA MGMT: DATA STORAGE :)

PERFORMANCE



DATA STORAGE, MIGRATION, LIFECYCLE MGMT

CHALLENGES IN DATA MGMT



DATA SILOSES, SERVICES FRAGMENTATION, DATA DUPLICATION

CHALLENGES IN DATA MGMT



DATA SILOSES, SERVICES FRAGMENTATION, DATA DUPLICATION

THE SOLUTIONS

PART ONE

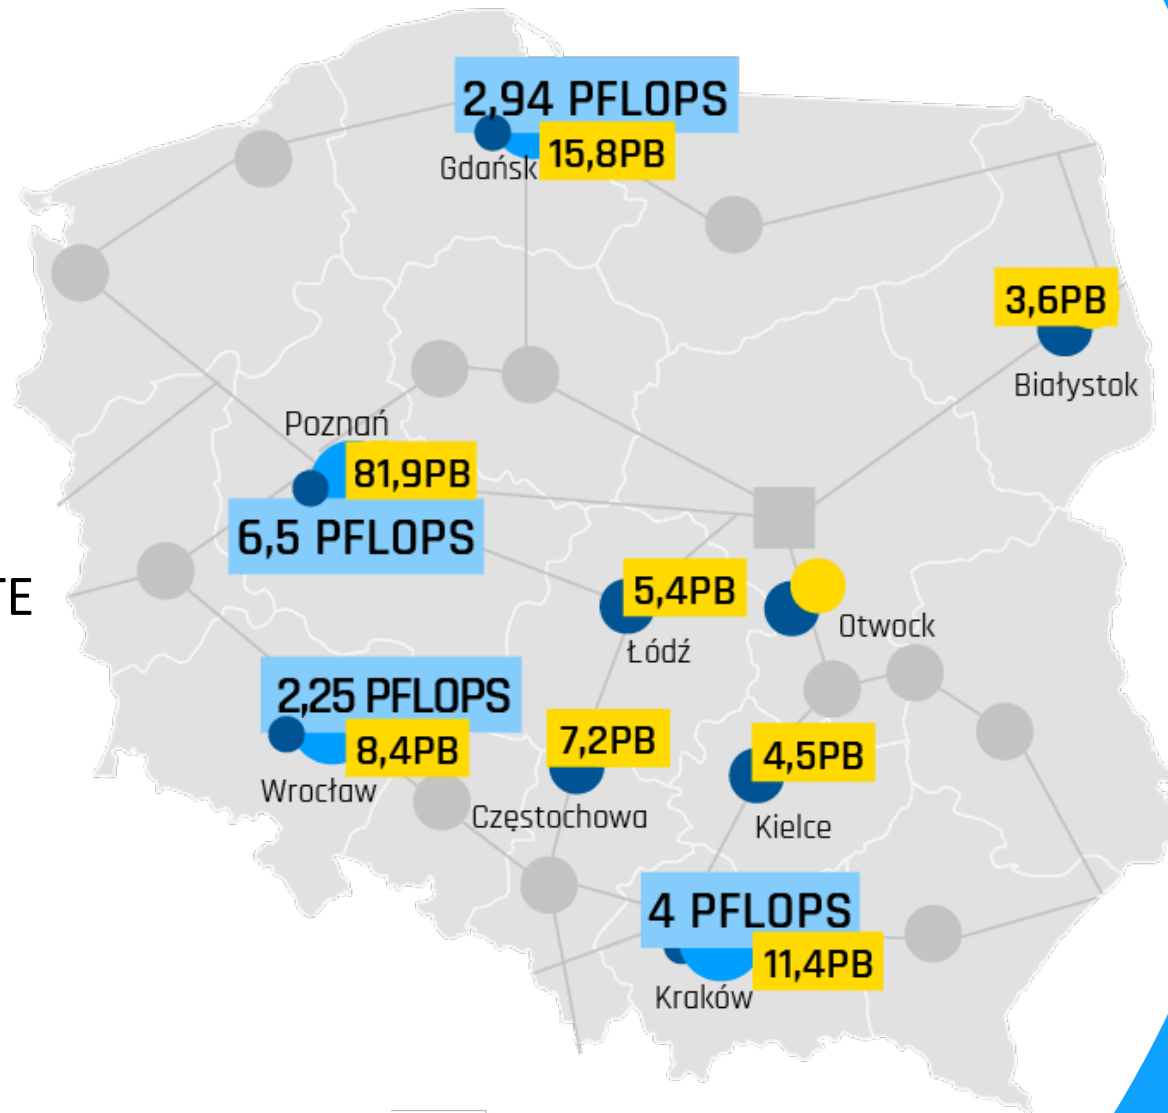




PRACE: collaboration in advanced
computing in Europe
(nationally funded, local implementation)

STORAGE SYSTEMS CO-LOCATED WITH COMPUTE (HPC, CLOUD)

- **100+PB CAPACITY**
IN HPC / CLOUD
COMPUTING CENTERS
- PROVIDED LOCALLY,
AND THROUGH
NREN NETWORK
- DISK + NVMe STORAGE
- CO-LOCATED WITH COMPUTE



DATA MGMT SERVICES

- THE 'CLASSICS'

- HPC STORAGE:
 - LUSTRE
 - NVME
- CLOUD STORAGE:
 - RBD ON CEPH
 - SSD TIERS
- ARCHIVE STORAGE:
 - CEPHFS
 - S3

OK, BUT THESE ARE THE BASICS

INFRASTRUCTURE, LOW-LEVEL SERVICES

HOW TO INTEGRATE?

DATA, STORAGE, COMPUTE ANALYSIS,
EXPLORATION, EXPLOITATION?

THE SOLUTION

PART 2





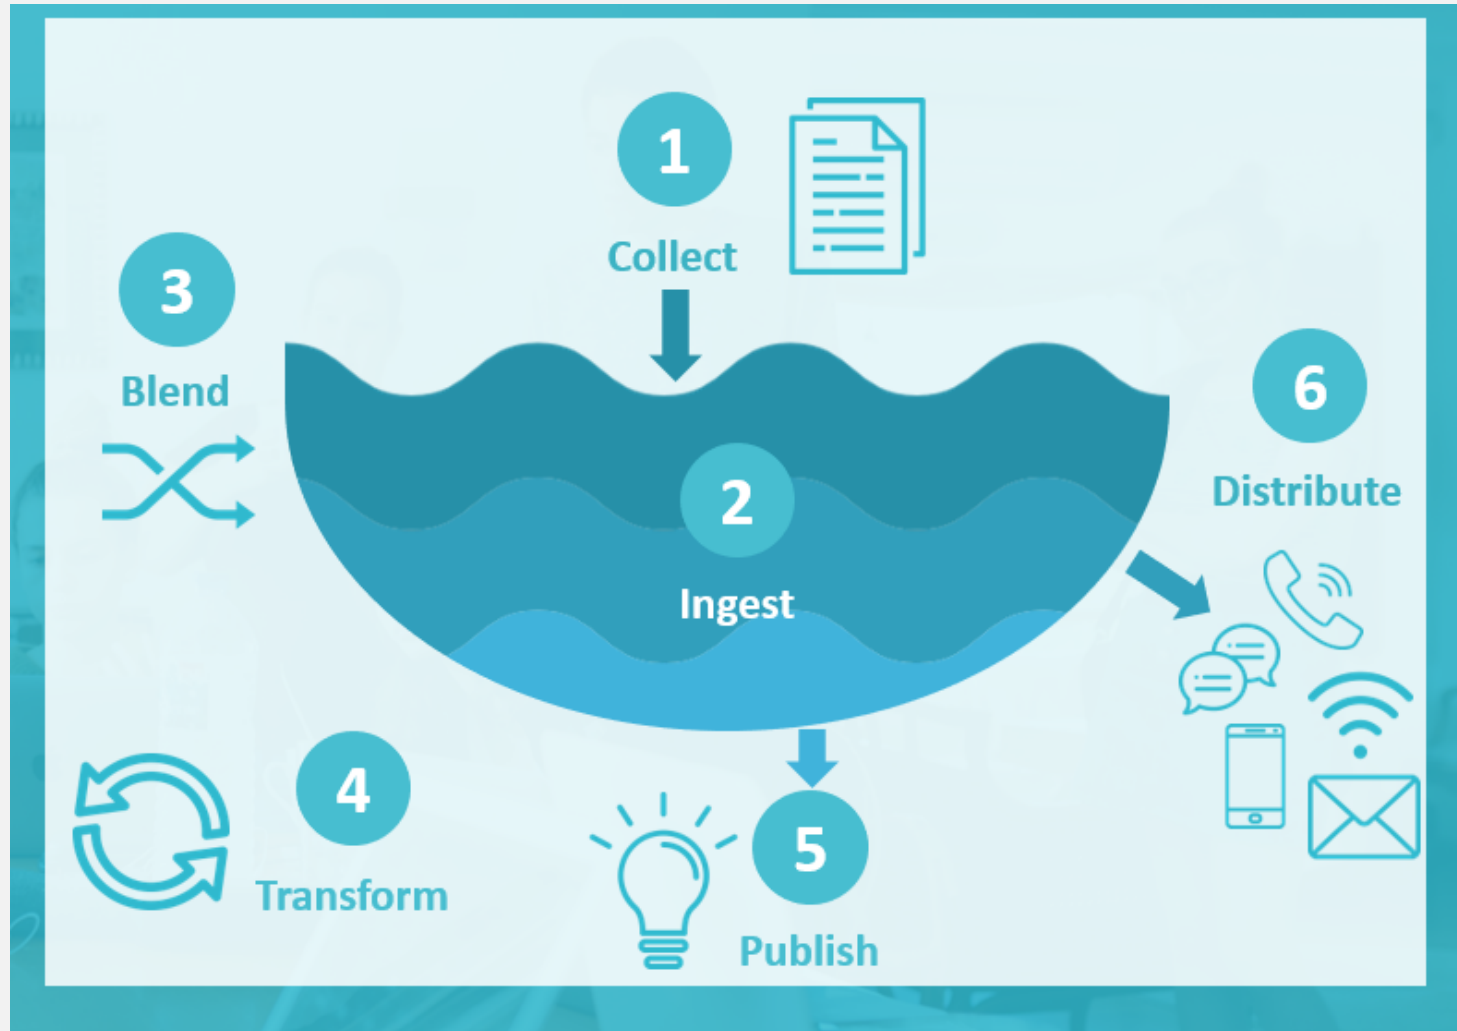
The infrastructure
for storage, access,
and processing of large data sets
in HPC, BigData i AI models
based on the data-lake concept



Project partners:




DATA LAKE TO THE RESCUE (I)



<https://www.ecloudvalley.com/what-is-datalake-and-datawarehouse/>

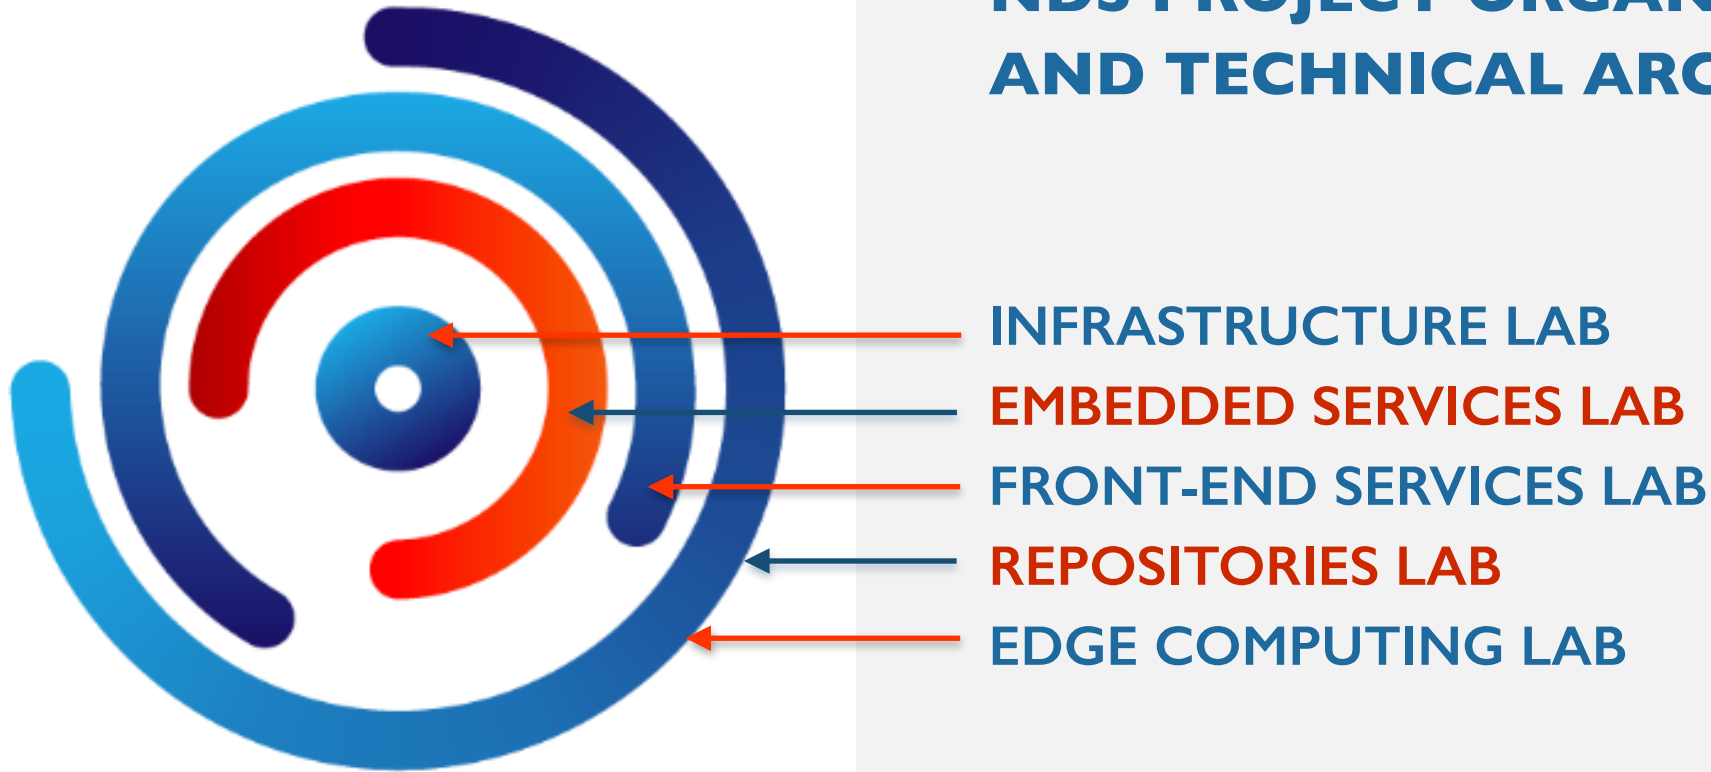
DATA LAKE TO THE RESCUE (2)

Data Lakes Features



<https://www.analyticsvidhya.com/blog/2022/10/top-data-lakes-interview-questions/>

NDS PROJECT ORGANISATION AND TECHNICAL ARCHITECTURE:



DATA LAKE IN THE **NDS** PROJECT



INFRASTRUCTURE LAB

ONLINE STORAGE:

- HPC STORAGE - SCRATCH
 - LUSTRE: **20 PBs** of HDD storage
 - NVME: **1PB** of NVMe cache (5%)
 - 500 GB/s, 1MLN IOPS
- HPC USERS' DATA STORAGE:
 - **44 PB** HDD storage
 - **2,2 PB** NVMe cache (5%)
 - 440GB/s, 2MLN IOPS
- "MASS" STORAGE:
 - DISK SERVERS: **50PB** of disk storage
 - NVMe SERVERS: **5PB** of NVMe storage
- ARCHIVE STORGE:
 - **100 PB** of disk storage
 - CEPH, S3

LUSTRE

ZFS

CEPH

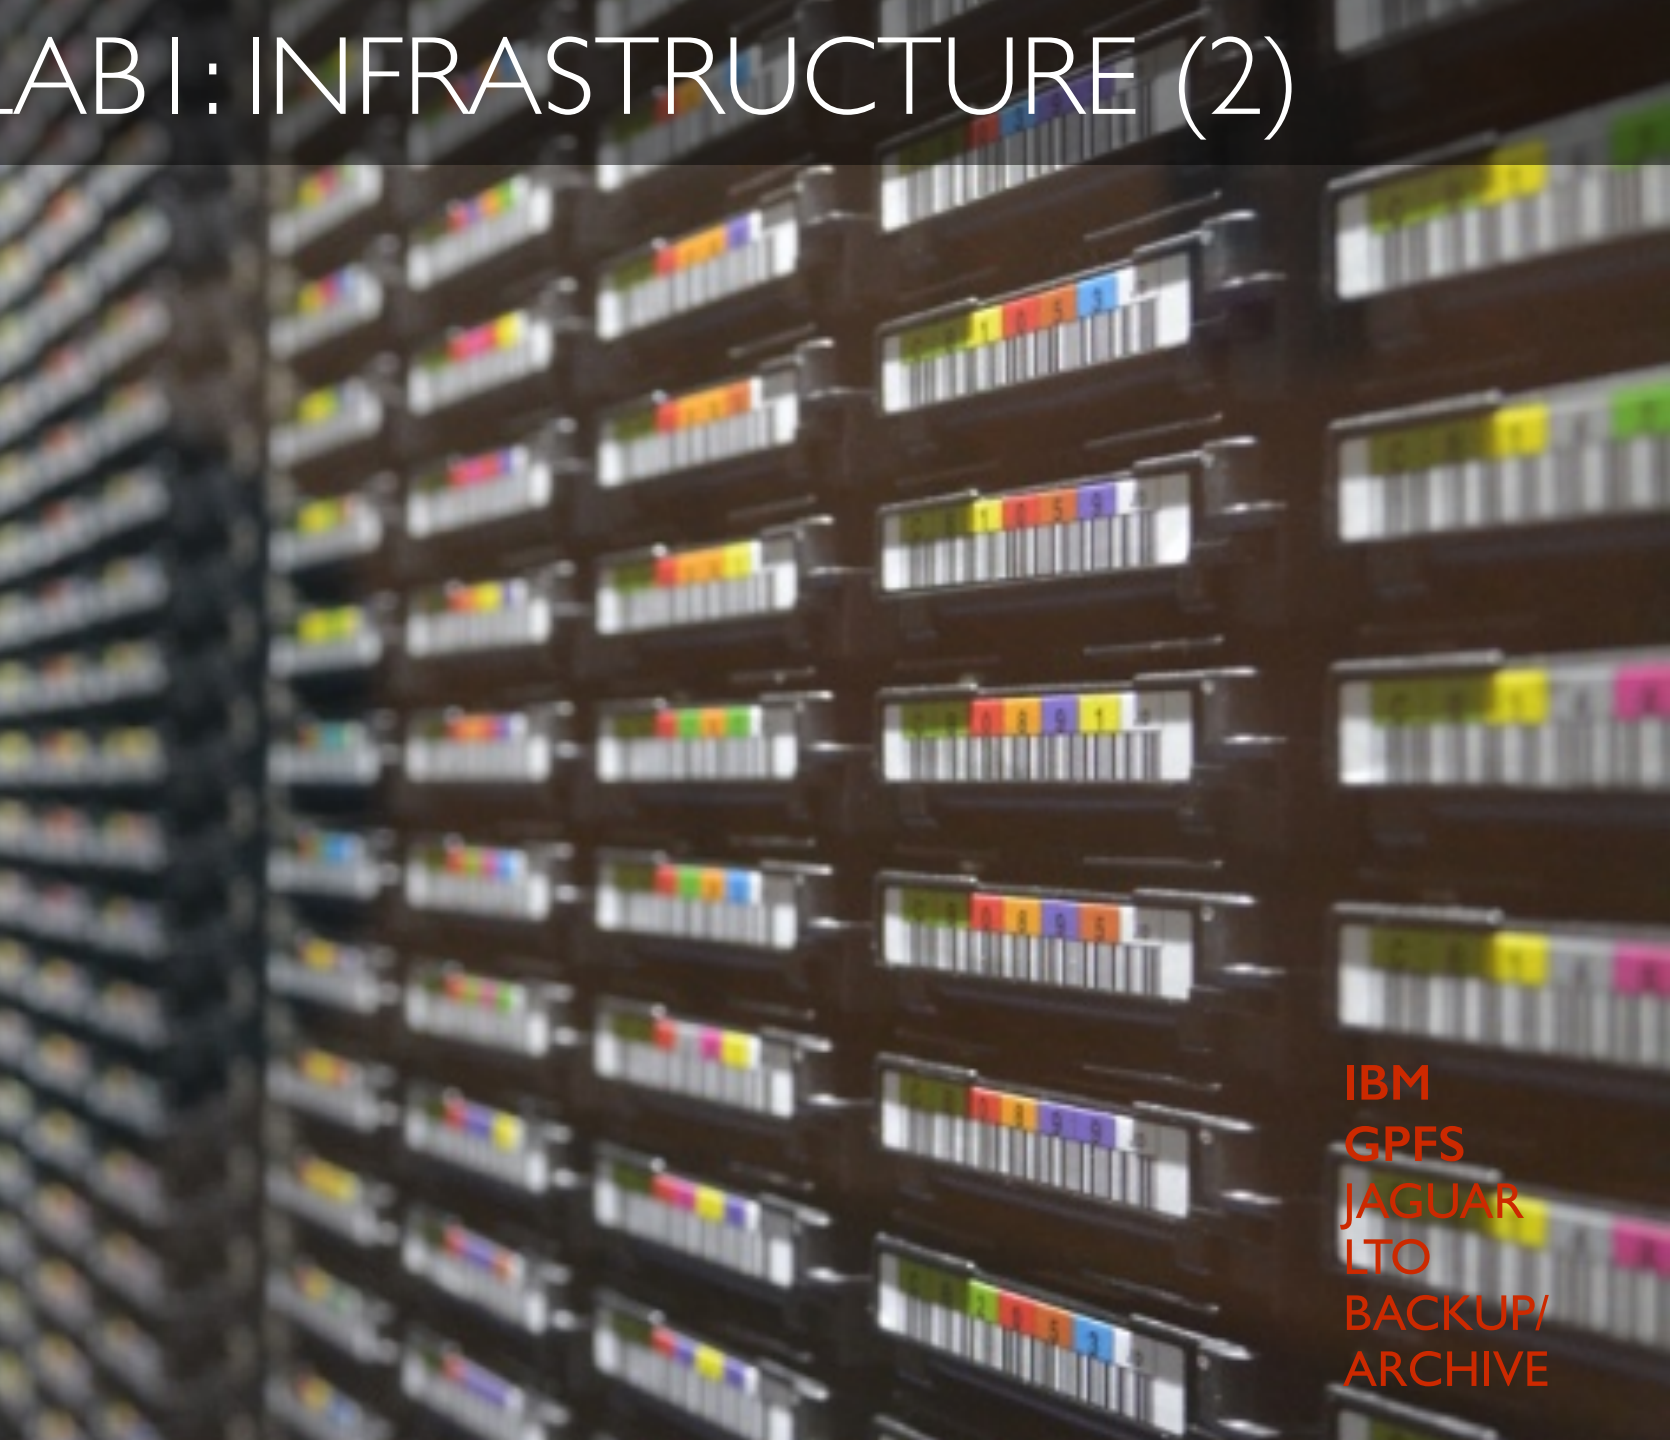
NVME

S3

AMAZON

OFFLINE STORAGE:

- TAPE SYSTEMS:
 - 500 PB IBM JAGUAR/LTO9
 - 5 PB HDD DISK CACHE (1%)
- HIERARCHICAL STORAGE (HSM):
 - FILE INTERFACE
 - S3 INTERFACE
- SERVICES: BACKUP/ARCHIVE
 - FOR DATA CENTRES
 - FOR END-USERS



IBM
GPFS
JAGUAR
LTO
BACKUP/
ARCHIVE

DATA LAKE IN THE **NDS** PROJECT



EMBEDDED SERVICES LAB

SERVICES:

- **BACKUP** as a SERVICE:
 - COMMERCIAL ENGINE BASED
 - OPENSOURCE ENGINE BASED
- LONG-TERM **ARCHIVE** as a SERVICE
 - HPC STORAGE -> ARCHIVE INTERFACE
 - S3 INTERFACE TO TAPE
 - DATA PACKAGING
 - ARCHIVE INDEXING AND SEARCH

BAREOS
STORWARE

LUSTRE/ROBIN HOOD

ZENKO CLOUD SERVER
MIN.IO / CEPH

SERVICES:

- **DOCUMENT DATABASES** as a service:
 - ELASTIC SEARCH
 - OPENSEARCH
- **CONTENT DISTRIBUTION** as a services:
 - KALTURA
 - WOWZA / UNIFIED STREAMING
- **DATA LAKE** AS A SERVICE:
 - DATA ACCESS W/ SEARCH & FILTERING
 - ON-THE-FLY DATA CONVERSION
 - AUTOMATED PROCESSING (LAMBDA_s)

KALTURA
WOWZA
UNIFIED STREAMING

DATA LAKE IN THE **NDS** PROJECT



FRONT-END SERVICES LAB

SERVICES:

- **DATA ACCESS SERVICES:**

- SYNC & SHARE
- FEDERATED DATA ACCESS
- DATA MIGRATION: IN/OUT SERVICE

- **DATA-CENTRIC APPLICATIONS:**

- DATA SCIENCE ENVIRONMENTS
- COLLABORATION APPLICATIONS
- DEVELOPERS' TOOLS

SEAFILE
NEXTCLOUD

ONE DATA

JUPYTER

ONLYOFFICE, COLLABORA
GITLAB

SERVICES:

- EXTRA SECURITY LAYER:
 - CRYPTOGRAPHY INTEGRATION
 - DATA ACCESS CONTROL
 - RECOMMENDATIONS, REGULATIONS...

ENCRYPTION

DATA LAKE IN THE **NDS** PROJECT



REPOSITORIES LAB

SERVICES:

- **DATA REPOSITORY AS A SERVICE:**

- PROPRIETARY SOFTWARE-BASED
 - DLIBRA
 - OTHER SOLUTIONS
- OPEN-SOURCE SOFTWARE-BASED
 - DATAVERSE
 - DSPACE

DLIBRA

DATAVERSE

DSpace

SERVICES:

• TYPES OF REPOSITORIES DELIVERED:

- INSTITUTIONAL REPOSITORIES
VS GENERIC / CATCH-ALL REPOS
- MULTI-DOMAIN REPOS
VS RESEARCH DOMAIN REPOS:
 - EARTH SCIENCE
 - BIOLOGY / MEDICINE
 - SATELLITE
 - PHYSICS

DATAVERSE

DSPACE

THE ACTUAL RESULTS

PART 3 OF THE SOLUTION



READY-TO-USE SERVICES:

- **DEPLOYED IN THE NDS INFRA**

- TAILORED TO A DOMAIN OR A USER GROUP
- GENERAL-AUDIENCE / CATCH ALL
- DEMONSTRATION SERVICES / PROTOTYPES

- **BLUEPRINTS:**

- DEPLOYMENT PLAYBOOKS
- DOCUMENTATION + CODE (INTEGRATION CODE, AUTOMATION SCRIPTS)



USE-CASES

PART 3 OF THE SOLUTION



USE-CASES

THREE LEVELS OF INTEGRATION

- **INFRASTRUCTURE (STORAGE) - RADIO-ASTRONOMY LONG-TERM ARCHIVE**

- STORAGE BACK-END FOR CACHING NODES (S3, CEPHFS)
- COUPLING WITH PROCESSING INFRASTRUCTURE

- **PLATFORM LEVEL (STORAGE + PaaS) - BIOLOGY REPOSITORY**

- IaaS: CPU + RAM FOR DATAVERSE APPLICATION
- STORAGE BACK-END FOR DATAVERSE: S3 FOR DATA OBJECTS
- PaaS: COMPONENTS FOR INDEXING & SEARCH

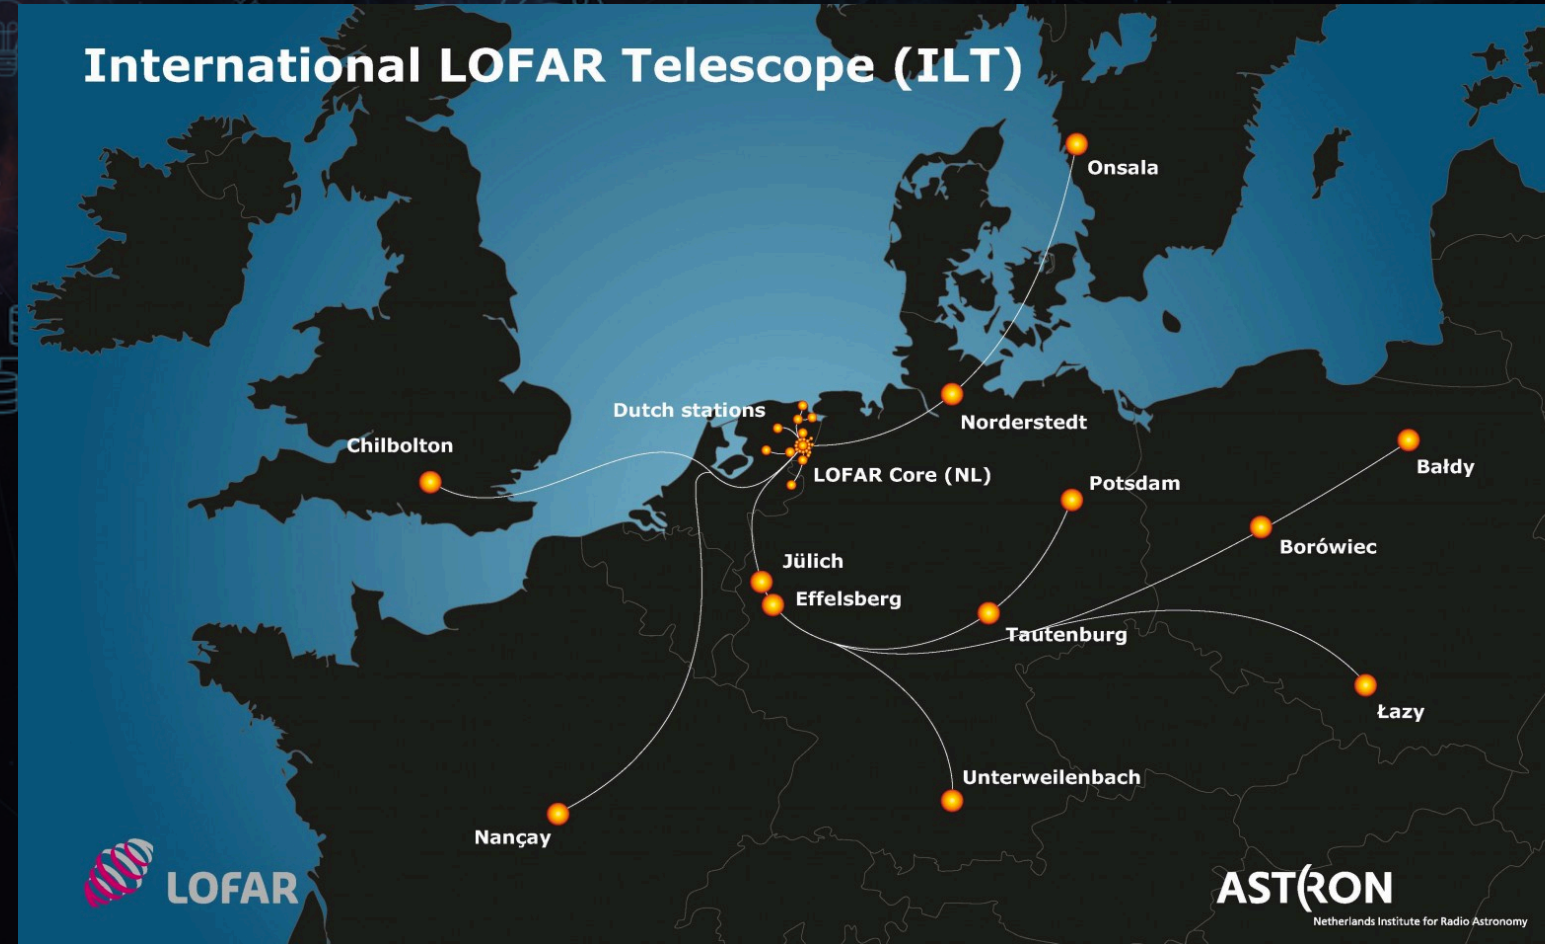
- **STORAGE + PROCESSING:**

- STORAGE: DATALAKE: DATA PERSISTENCE + ONLINE AVAILABILITY + QUERY-ABILITY
- PaaS: COMPONENTS FOR INDEXING & SEARCH
- COUPLING WITH PROCESSING INFRASTRUCTURE: CLOUD AND HPC CLUSTER

RADIO-ASTRONOMY USE-CASE

LOFAR LONG-TERM ARCHIVE:

- Acquisition, analysis, processing
 - Distributed data sources
 - Centralised correlation,
 - Distributed storage
 - Processing in HPC centres
 - Data access to project members
- LTA (Long-Term Archive) in PSNC:
 - Constant data stream: **10 Gb/s**
 - Since 2015, now **20PB+** of data
 - Tape storage + disk cache




**KEEP
CALM
AND
TAPE**



http://www.astron.nl/dailyimage/pictures/20160518/First_byte_Poznan.png

RADIO-ASTRONOMY USE-CASE

- LONG-TERM ARCHIVE
- DATASET: 20+ PB
- RESEARCHERS NEED:
 - Long-term access to astronomical data need to keep reference historical data
 - Scheduled and ad-hoc processing for detecting the astronomical events
 - Ad-hoc data access
 - high-performance download

LOFAR USE-CASE DISCUSSION

- CHALLENGES:
 - Large data volume, pressure on performance
 - Custom software stack (dCache, GridFTP)
- NDS system benefits:
 - Cost-efficient data infrastructure
 - commodity disk servers
 - Distributed storage Ceph
 - Distributed, multi-node dCache
 - I/O parallelism - no bottle-necks
 - Coupling with cloud and HPC:
 - Co-located data analysis, e.g. for content re-coding



BIOLOGY USE-CASE

- LARGE-SCALE DIGITISATION OF SPECIMENS: PLANTS, ANIMALS, FUNGHI
- PARTNER: ADAM MICKIEWICZ UNIVERSITY FACULTY OF BIOLOGY
- RESEARCHERS NEED:
 - Collaboration, ease of use, document editing
 - Dropbox-like experience
 - Persistent data storage in a multi-tenant repository with indexing and search
- DATASET:
 - 100+ TB,
 - 100s 1000s of objects



BIOLOGY USE-CASE

- **SHORT TERM REPOSITORY AND COLLABORATION SPACE:**

- LARGE VOLUMES: 10-100s TBs
- EASE OF USE/ACCESS
- FUNCTIONALITY
- IMPLEMENTATION:
 - Seafile, NextCloud-based sync & share as a service



- **(OPEN) DATA REPOSITORIES:**

- LARGE DATA VOLUMES, NUMBER OF OBJECTS
- SUPPORT FOR META-DATA, INDEXING, SEARCH
- PRESENTATION OF VARIOUS DATA FORMATS:
 - TEXT, DOCUMENTS
 - GRAPHS, 3D
- IMPLEMENTATION:
 - DataVerse as a service
 - Alternatively - DataCite as a service



BIOLOGY USE-CASE DISCUSSION

- CHALLENGE:
 - Different interactions with the data sets, in various stages of digitisation
 - First - ease of use, interaction
 - Then - storage persistency, query ability, efficient access
- NDS system benefits:
 - Integrated sync & share provided the dropbox-like experience
 - Repository software enabled multi-tenancy (spaces) and flexible access control (1st internal repository, then public)
 - Platform-provided indexing and search engine: OpenSearch enables data exploitation and exploration
 - Sync & share -> data repository deposit mechanism automatised migration

SECURITY USE-CASE ENRICHED MALWARE DATA LAKE (EMD)

- **IDEA**

- Automated malware gathering
- Defined unified analysis workflows @HPC
- Multidimensional queries for results
- Threat hunting - both live and retro

SECURITY USE-CASE ENRICHED MALWARE DATA LAKE (EMD)

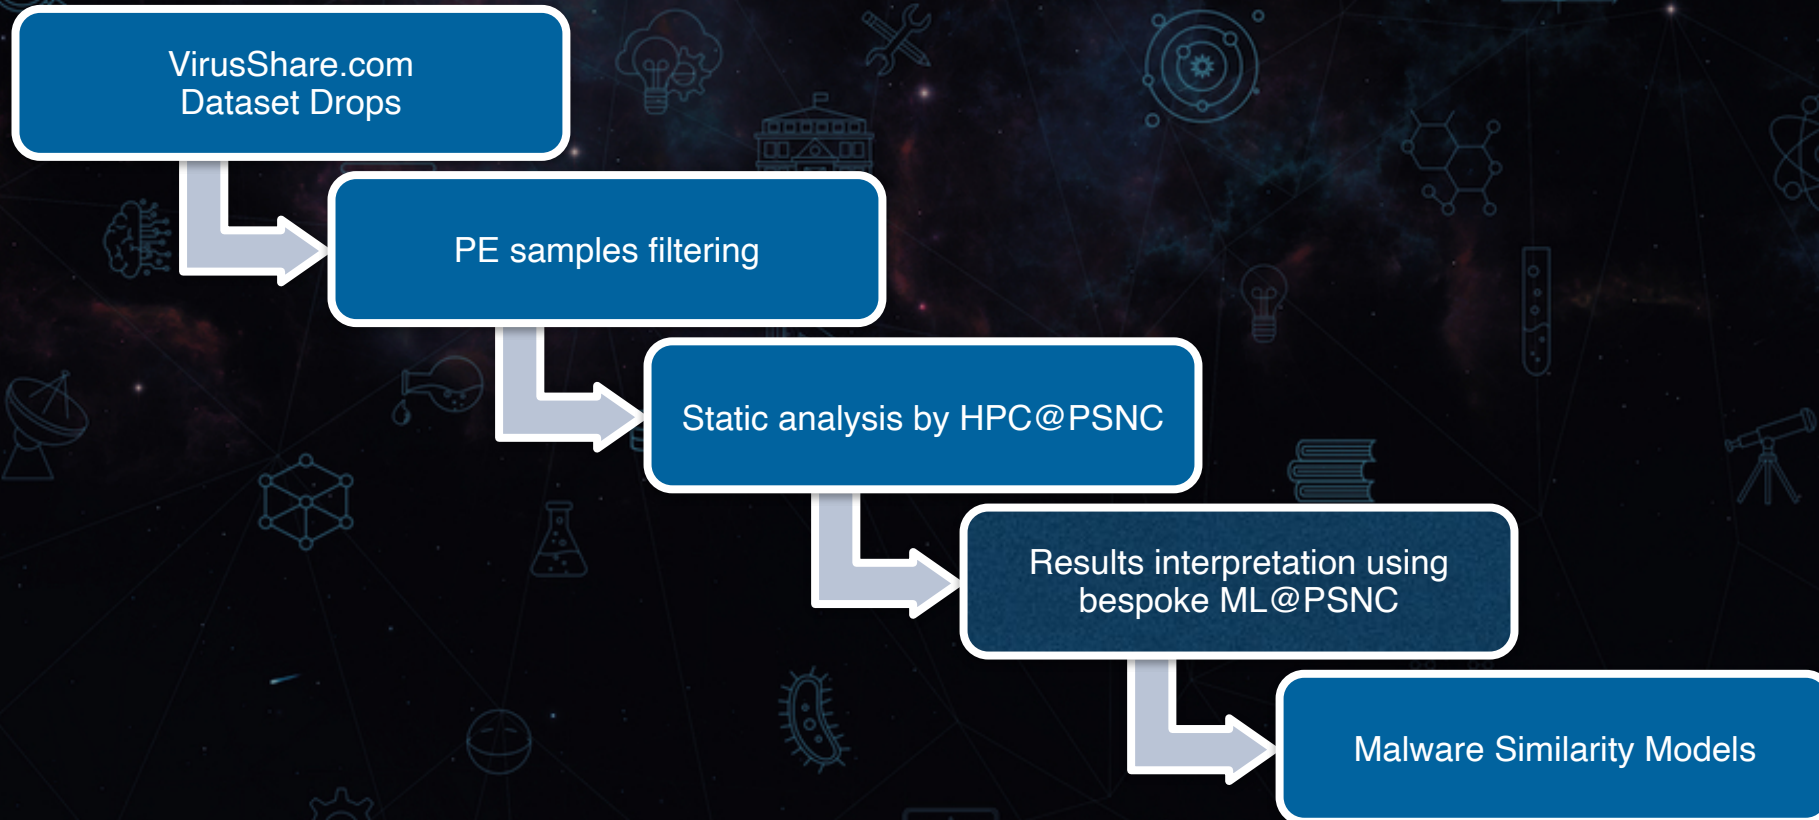
- **STORED DATA**

- Raw malware samples
- Analysis results serving as rich „meta” data
 - Static analysis
 - Sandbox results (incl. RAM, disk & network dumps)

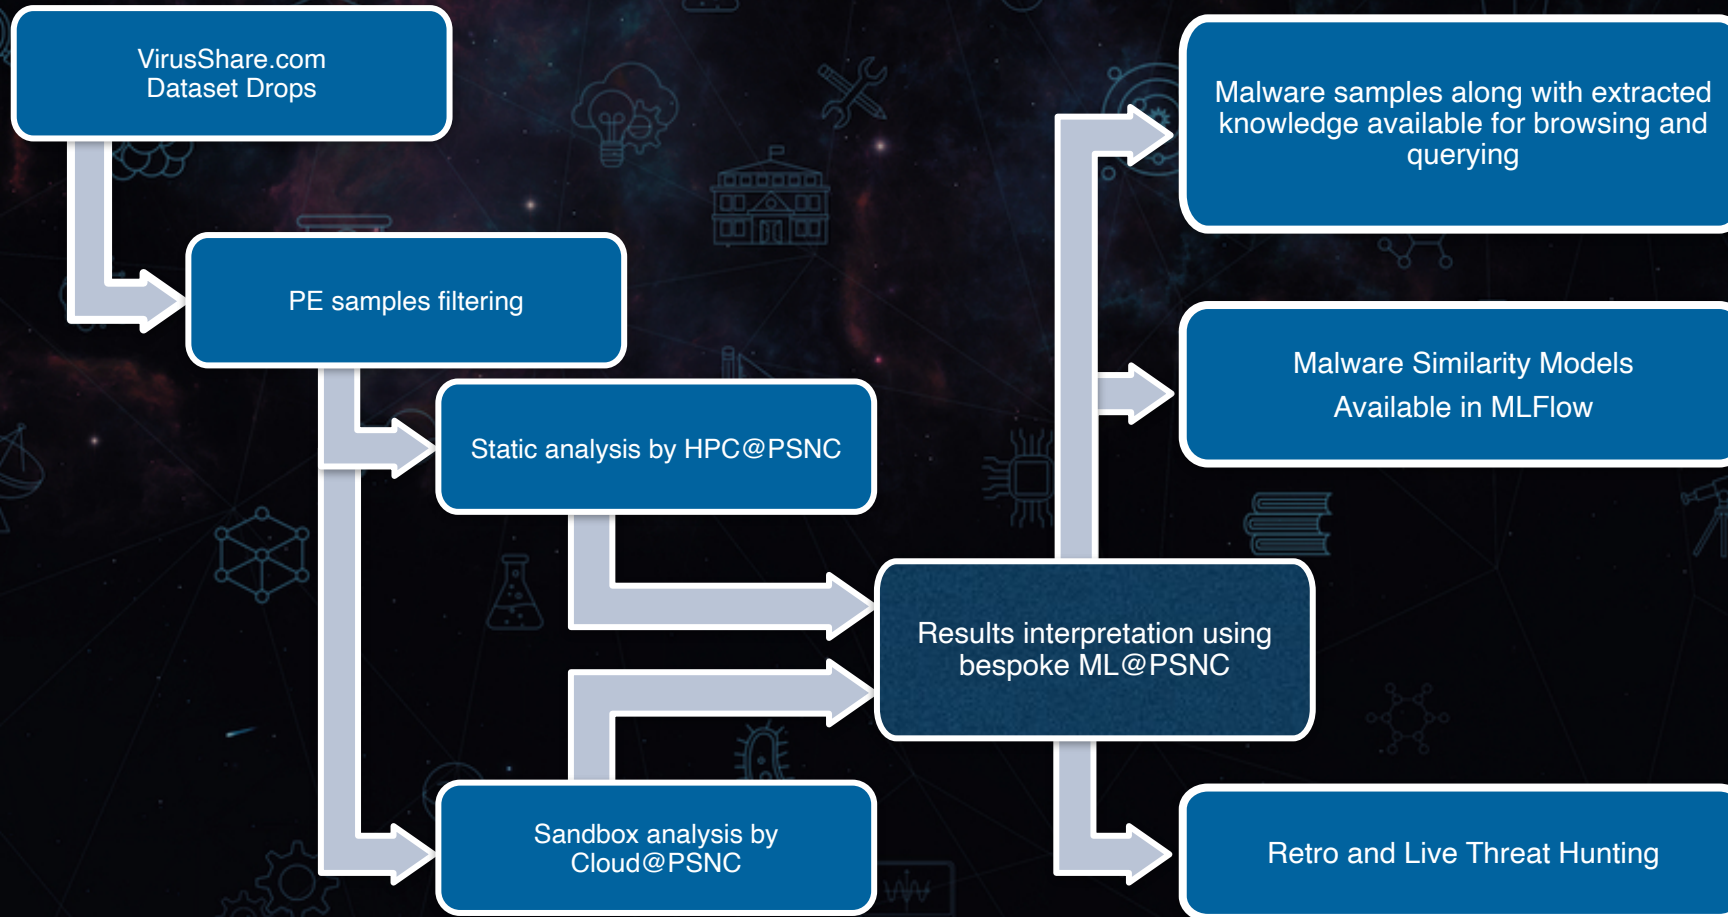
- **NUMBERS** (so far)

- **17 TB** (~20 mln) raw malware samples (Windows PE)
- **250 TB** decompiled sources
- **3 TB** malware capabilities metadata used for indexing

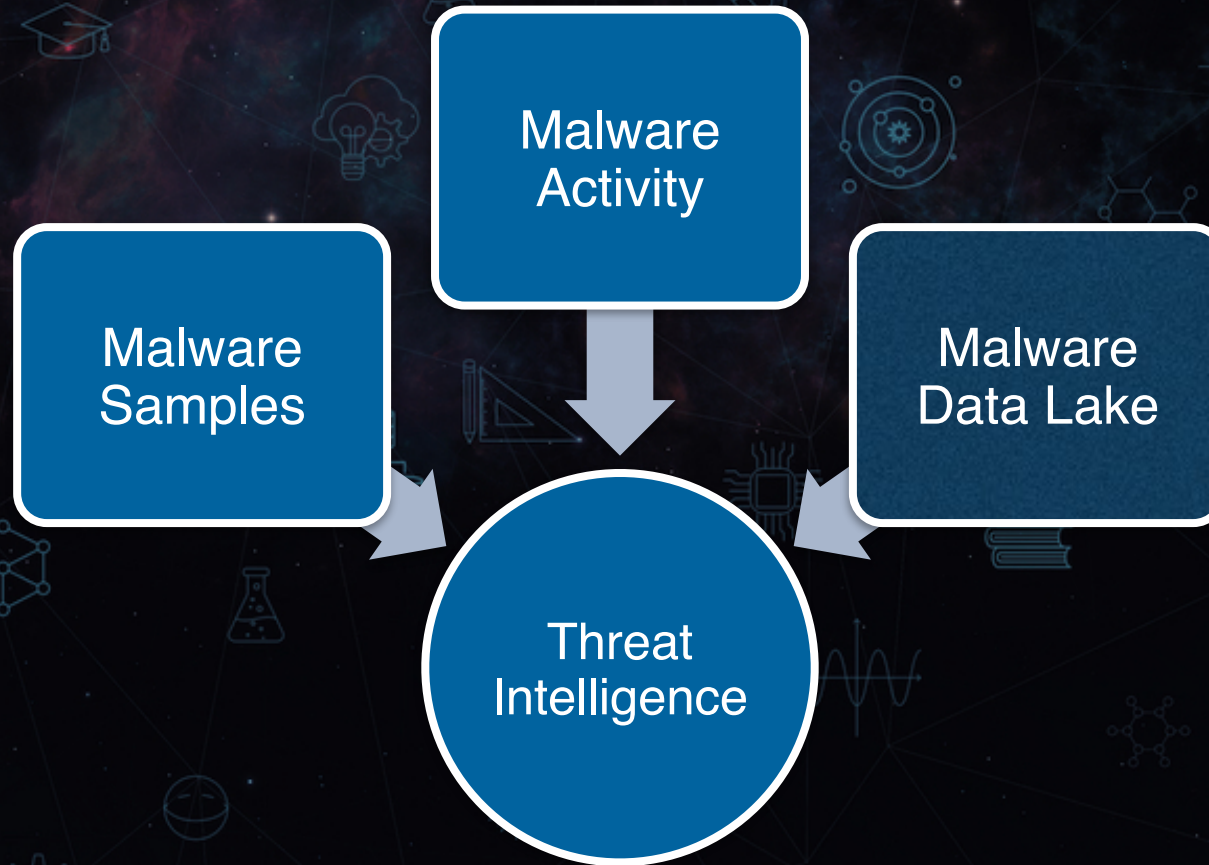
EMD POC FLOW



EMD FLOW




EMD WHY

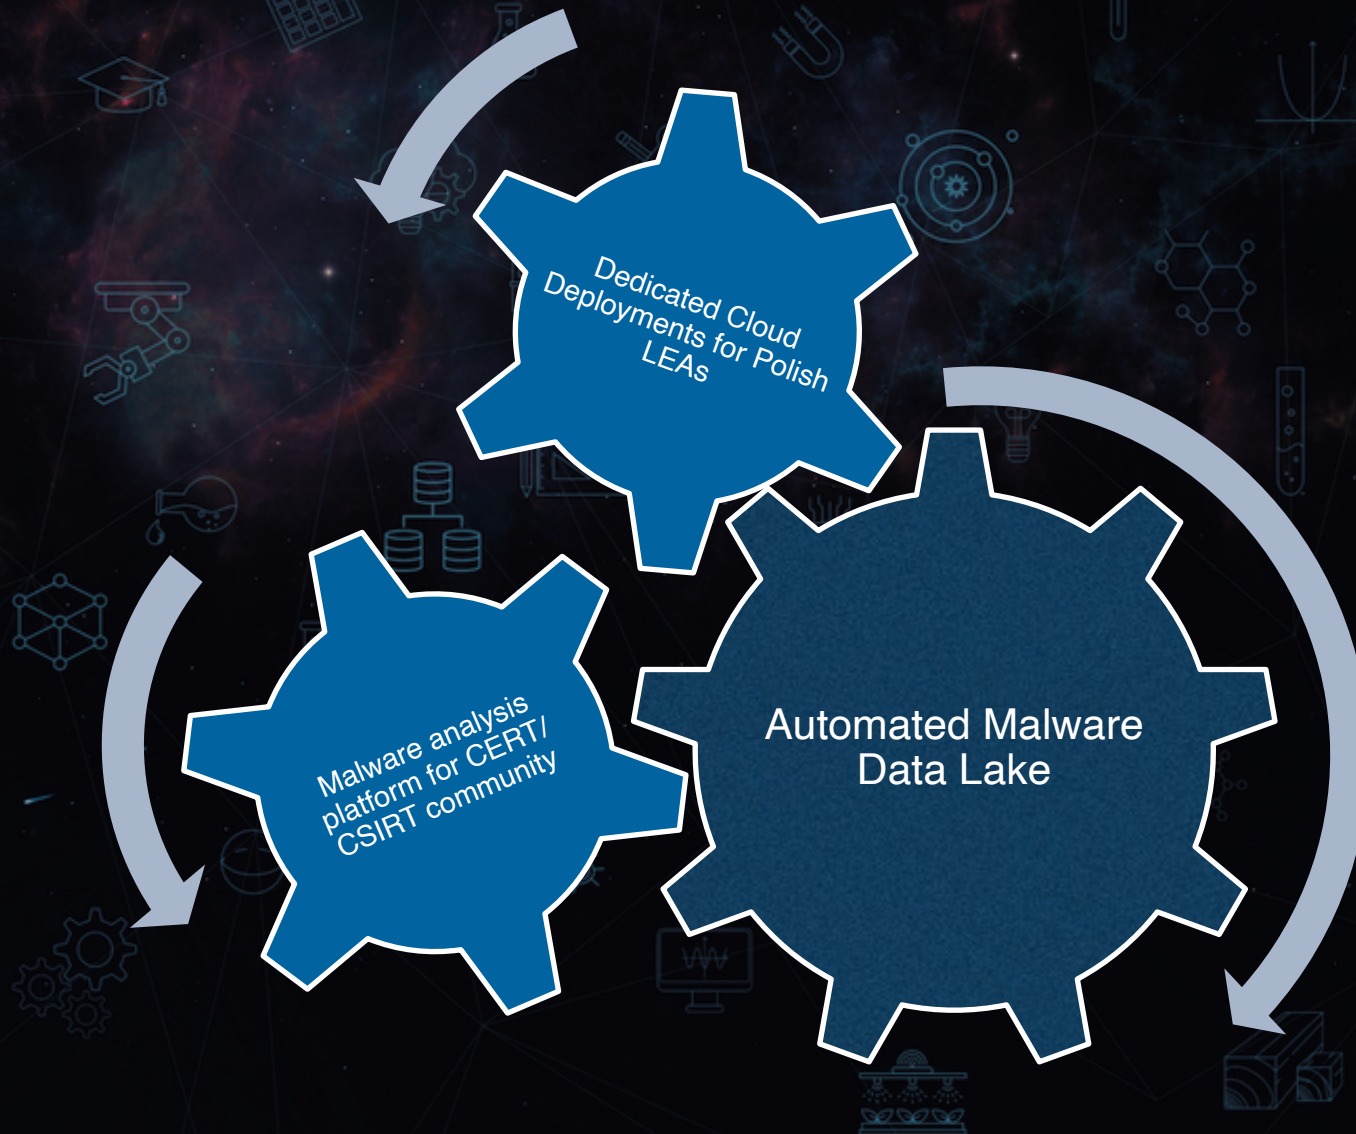


EMD WHY

- **DATA LAKE**
 - Automated malware gathering
 - Defined unified analysis workflows @HPC
 - Multidimensional queries for results
 - Threat hunting - both live and retro
- **(Planned) CSIRT SaaS**
 - Malware Analysis
 - Forensics support
- **(Intended) LEAs**
 - Dedicated projects
 - Bespoke extensions



EMD SYMBIOSIS



SECURITY USE-CASE DISCUSSION

- CHALLENGE:
 - Large dataset (Petabyte-scale)
 - Query-ability!
 - Can't go offline!
- NDS system benefits:
 - storage infrastructure enables data availability
 - platform-provided indexing and search engine: OpenSearch
 - tight coupling of HPC and Cloud computing clusters ensures low-latency, high-throughput data access

SECURITY USE-CASE DISCUSSION

BENEFITS

- Automated processing close to data:
- High performance, lower resource consumption
- Direct data access - simpler system design
- Simplified architecture
 - Services embedded in the data infrastructure
 - 'Server-less' architecture

SUMMARY - **NATIONAL DATA STORAGE**

AN INCARNATION OF DATA LAKE CONCEPT

- POSITIONING:
 - AMBITIOUS PROJECT WITH PERSPECTIVE OF A LONG-TERM IMPACT AND BACKING OF USERS
 - DATALAKE TO COLLECT, STORE, PROTECT AND ENSURE USAGE OF THE DATA (IN-LINE WITH OPEN DATA INITIATIVES)
- RESULTS:
 - TOOLS + BASIC SERVICES + BACK-END INFRASTRUCTURE
 - READY TO USE SOLUTIONS: DEDICATED, GENERIC
 - BLUEPRINTS
 - RE-DEPLOYABLE SOLUTIONS
- 'ARCHITECTURE':
 - MINIMUM DEVELOPMENT EFFORT
 - MOSTLY INTEGRATION OF EXISTING COMPONENTS
 - LONG-TERM SUSTAINABILITY OF RESULTS

NATIONAL DATA STORAGE
DATA-LAKE-LIKE
INTEGRATED
DATA MANAGEMENT
SERVICES FOR R&E
THANK YOU!!!

MORE INFORMATION:

[HTTPS://KMD.PIONIER.NET.PL/](https://kmd.pionier.net.pl/)



Data mgmt team:
Krzysztof Wadówka,
Adam Prycki,
Eugeniusz Pokora,
Filip Blicharczyk,
Jan Bróździak,
Krzysztof Błoniarz,
Maciej Brzeźniak

Collaboration (security team):
Mikołaj Dobski (security)
Gerard Frankowski (security)

Collaboraiton (cloud, net, DC):
Jacek Kochan (cloud)
Sebastian Petruczynik (net)
Rafał Mikołajczak (DC)

Director of Data Processing
Technologies Division:
Dr Norbert Meyer