



# IBM at HPC User Forum 4/19/23

## Attendees

- **Mike Dye, US Federal Technology, DOE**
- **Tim McGinnis, US Federal Technology, POWER Systems, DOE, NASA**
- **Jim Bonfils, US Federal Technology, Storage, DOE**
- **John Unthank, US Federal Technology, Storage, DOE, NASA, NOAA**
  
- **Robert Eades, IBM Quantum Computing, Govt Science Program**
  
  
- **IBM Quantum / Active Archive / IBM POWER Systems**

# Development Roadmap

Executed by IBM   
On target 

## IBM Quantum Computing Summit November 2022


[www.ibm.com/quantum/roadmap](http://www.ibm.com/quantum/roadmap)

2019                      2020                      2021                      2022                      2023                      2024                      2025                      Beyond 2026


Model  
Developers



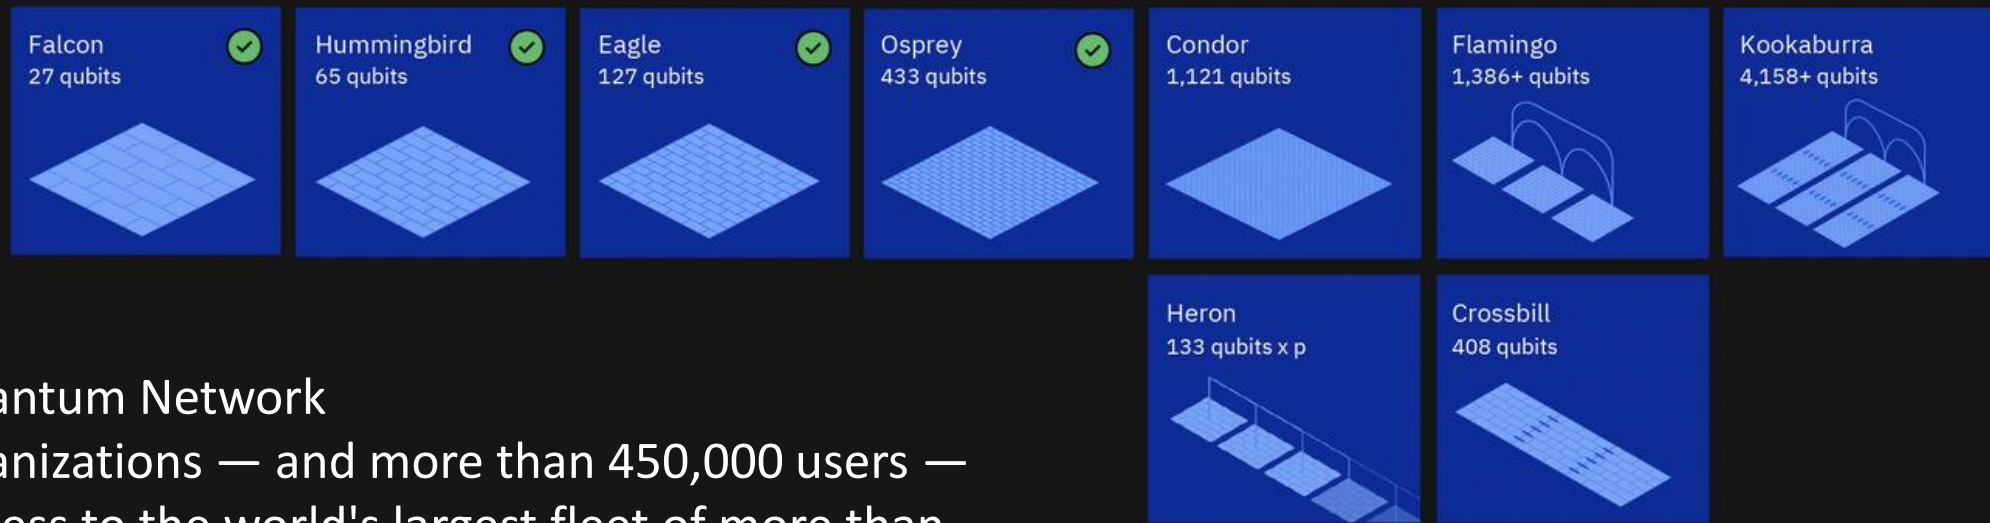
Algorithm  
Developers



Kernel  
Developers



System  
Modularity



### IBM Quantum Network

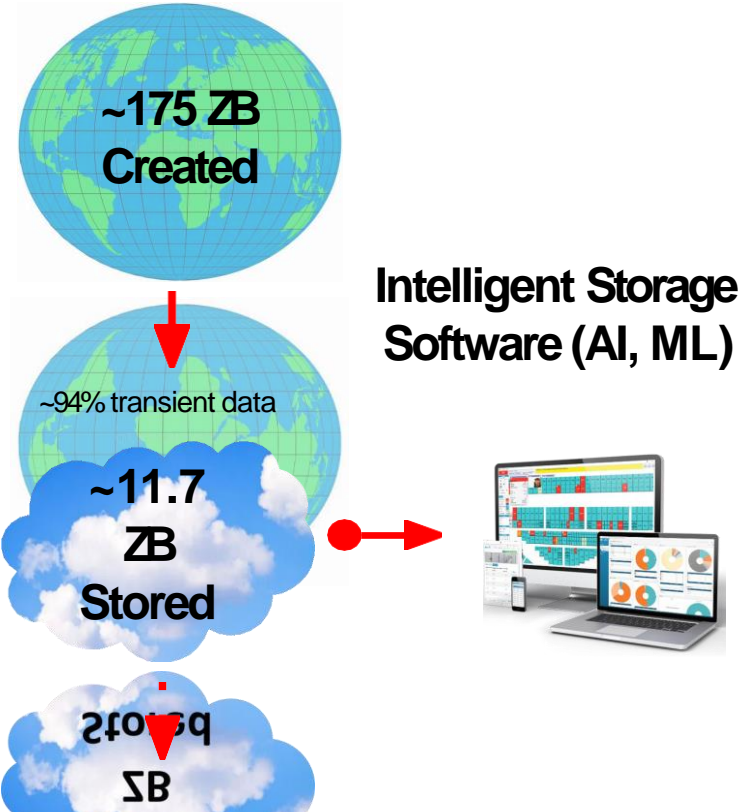
200 organizations — and more than 450,000 users —  
with access to the world's largest fleet of more than  
20 quantum computers accessible in the cloud.

# The Tiered Storage Model

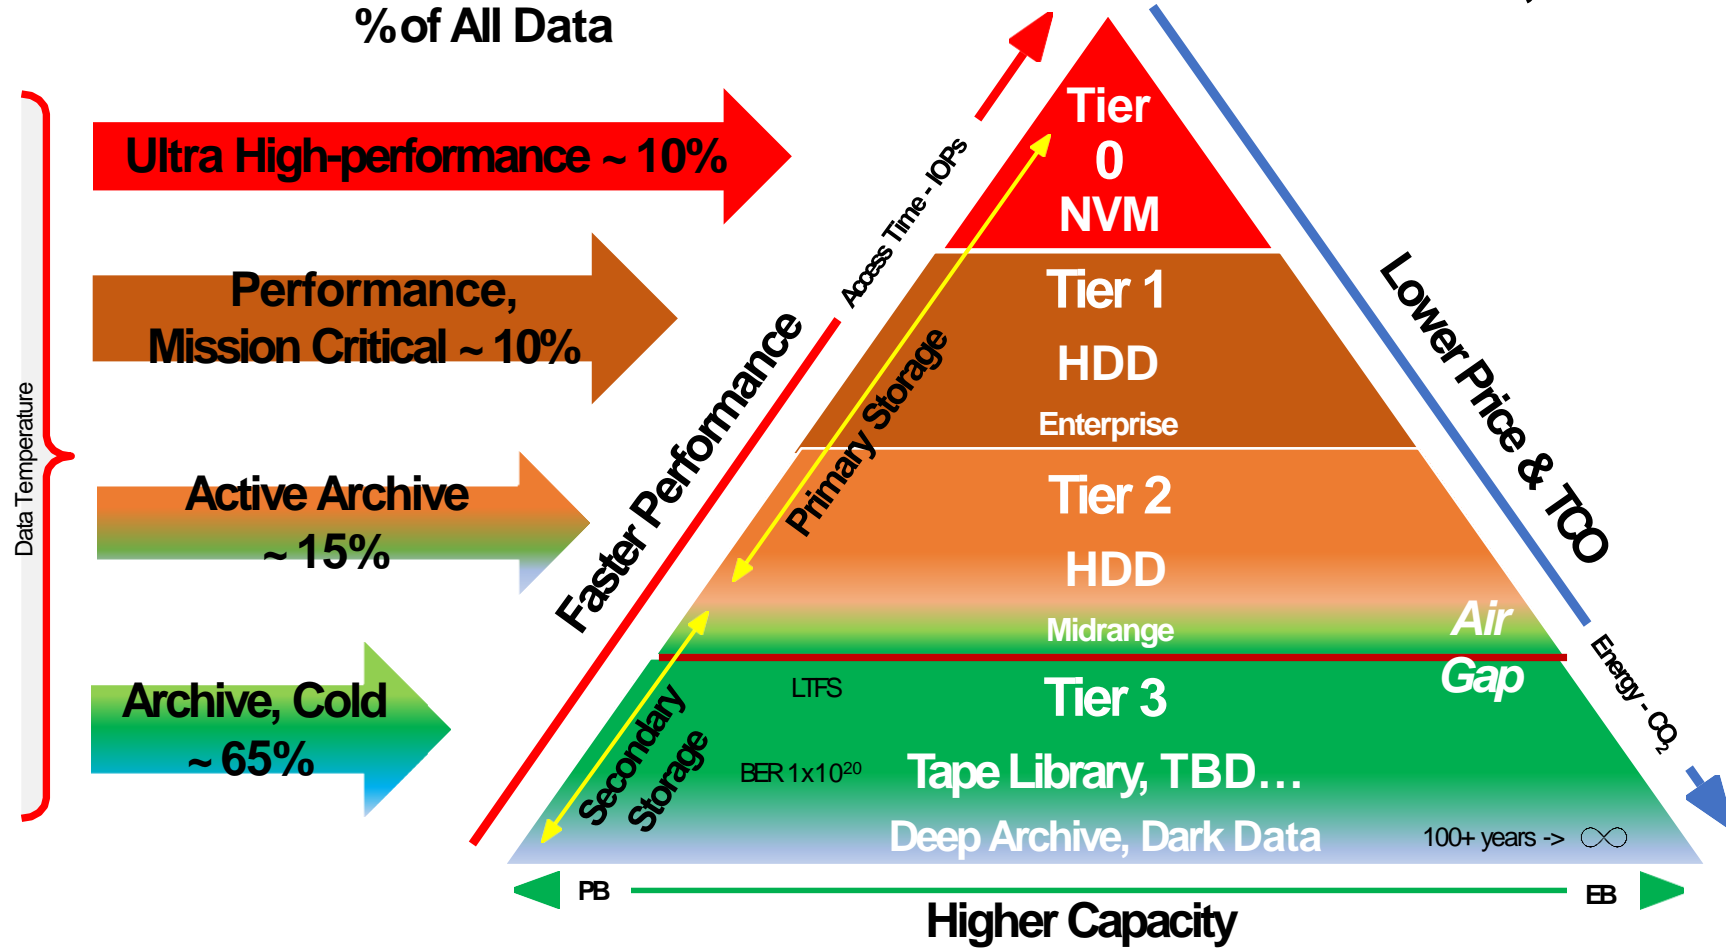
## Zettabyte Era Pushes Limits for Storage Technology

Source:  
Fred Moore,  
Horison, Inc.

By 2025  
Source: IDC



Optimized Data Classification by Tier  
% of All Data



Legacy, Unknown, Orphaned Data

IBM Future of Tape Briefing – Tucson morning of Sept 8 following Sept HPC User Forum

Youtube – IBM Diamondback Tape Library

# IBM POWER10 Systems

B  
Y  
T  
E  
S

## 2x Bytes from all sources (OMI, L3, L2, L1 caches\*)

- **4x 32B loads, 2x 32B stores per SMT8 core**
  - New ISA
  - Thread max 2x 32B loads, 1x 32B store
- **OMI Memory to one core**
  - 256 GB/s peak, 120 GB/s sustained
  - With 3x L3 prefetch and memory prefetch extensions

F  
L  
O  
P

## 2x Bandwidth matched SIMD\*

- **8 independent SIMD engines per SMT8 core**
  - Fixed, float, permute

## 4-32x Matrix Math Acceleration (vs POWER9)

- **4x 512b engines per core = 2048b results / cycle**
  - Matrix math outer products:  $A \leftarrow \{\pm\}A \{\pm\}XY^T$
  - Double, single, reduced precision

## IBM POWER10 E1080

- Up to 16 processors, Up to 240 cores
- Up to 64 TB memory
- No GPU

## IBM POWER10 E1050

- 4U Up to 4 processors , Up to 96 cores
- Up to 16 TB memory
- No GPU

- RHEL, SUSE, Ubuntu, and yes AIX

## Performance Examples / Use cases

- ✓ World record 4-socket SPEC CPU 2017 benchmark
- ✓ 2.2X greater efficiency per core
- ✓ 3x the throughput for heavy simulation i.e. LS DYNA
- ✓ 5X AI Inferencing improvement per socket Acceleration  
Based on IBM testing using ResNet-50, Pytorch, OpenBLAS on the same BERT Large with SqUAD v1.1 data set.
- ✓ High performance database