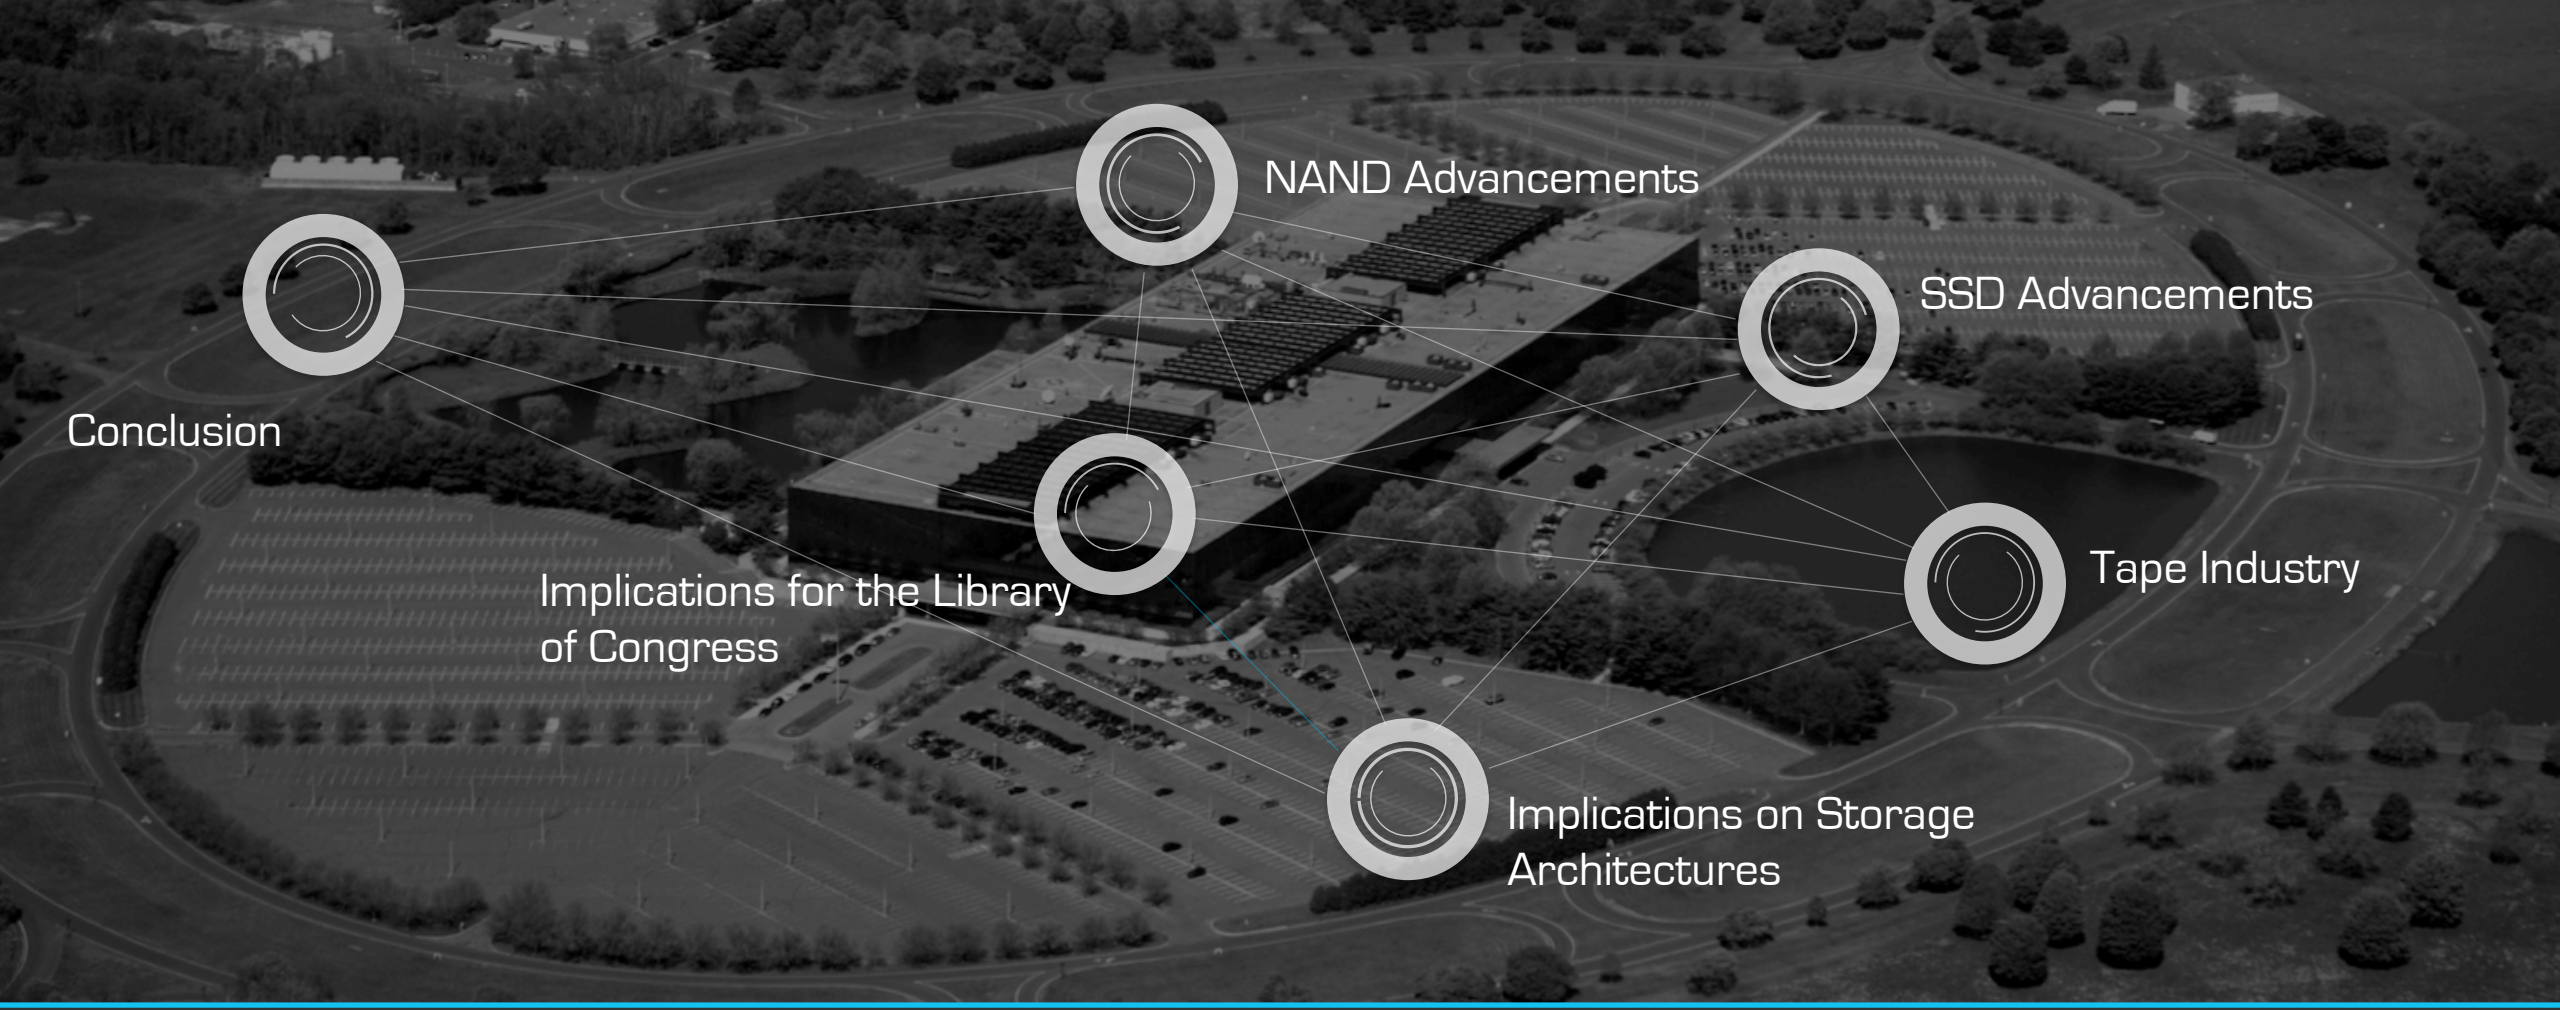


# Designing Storage Architectures 2017

Rob Peglar, SVP/CTO, Symbolic IO
rpeglar@symbolicio.com @peglarr

# NAND Advancements

Density is King – Will Approach $.05/GB to manufacture ($50/TB)

## 3D NAND SCALING

### Vertical Scaling


*Multiple stack of bilayers*

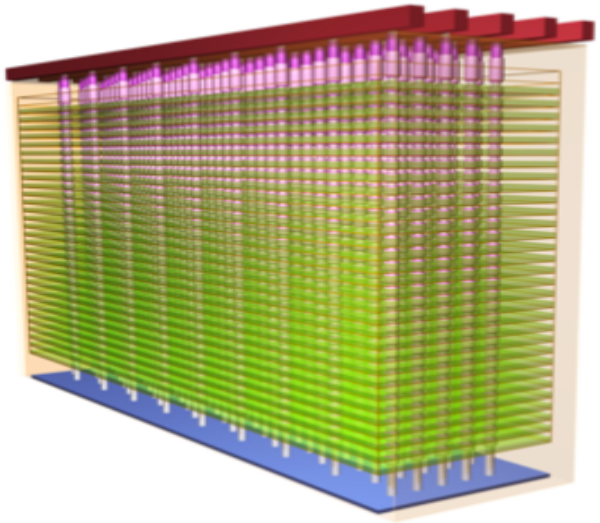

### Device scaling

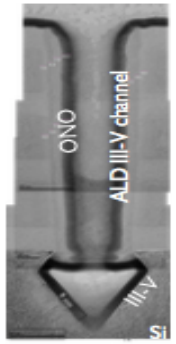


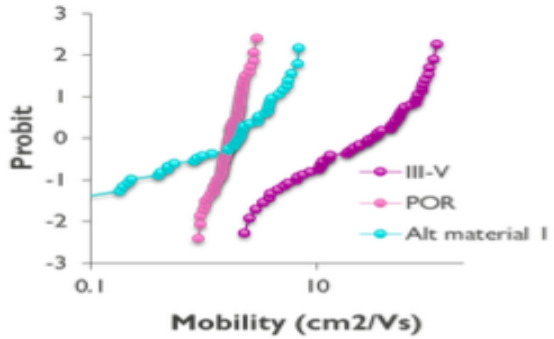


*Improved cell performance/mobility with III-V channel*

| Year | 2016-2017 | | 2018-2019 | | 2020-2021 | 2022-2023 |
|---|---|---|---|---|---|---|
| **Generation 3D** | L48 | L64 | L96 | L128 | L256 | 512 |
| **Die size (3b/cell)** | 256-512 Gb | 512Gb – 1Tb | 512Gb-2Tb | 1-3Tb | 2-6 Tb | 4-12Tb |
| **Hole CD** | 65-100 | 65-100 | 65-100 | 65-100 | 65-100 | 65-100 |
| **Slit pitch (# holes)** | 4 | 4 | 4-8 | 8 | 8 | 8 |
| **Vertical pitch** | 50-70nm | 40-60 | 40-60 | 40-50 | 40-50 | 40-50 |
| **BL CD** | 20 | 20 | 20 - 40 | ~40 | ~40 | ~40 |
| **Multiple stacks** | No | No | No | No | Yes (2-4) | Yes (4-8) |

Source: Intel, Projected Private Cloud CPU Sales

# SSD Advancements

## Density is Also King – in new Form Factors


- Very small SSDs – but very dense
- Current NAND density = 500 GB
- Will reach 4 TB by 2021

- Extremely large SSDs – in non-standard form factors
- SFF TA-1002 connector (pending standardization)
- Current NAND density = 1 PB (raw)
- Will reach 8 PB by 2021

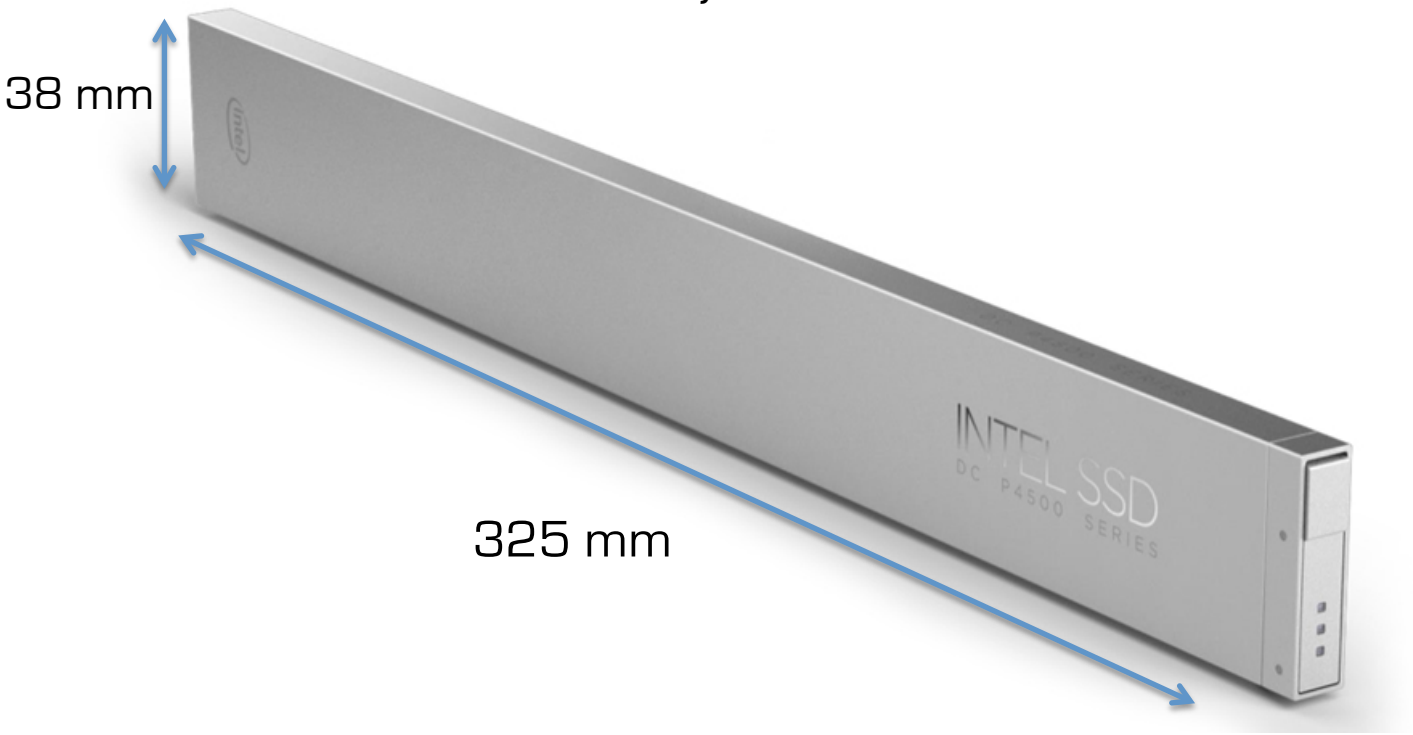


Source: Intel, Projected Private Cloud CPU Sales

# The Tape Industry

## Headed towards a Monopoly


- IBM and Spectra now the only manufacturers of robotic tape libraries
- Oracle/STK cancelled next-gen T10000 drive in March 2017
- One more mid-range Oracle/STK library release – then no more
- Everyone using T10000 drives/carts must migrate to LTO
- Spectra Logic emphasizing Black Pearl – tape-to-cloud, tape-to-flash
- TCO of operating large tape-based archives is bad and getting worse
  - Compared to dense NAND – and doesn't count cost of migrations
- It's not the cost of the tape cartridge – it's everything around it
  - Library & drive h/w, s/w, maintenance, support
  - Drive h/w, f/w, maintenance, support
  - Disk array(s), data channels, filesystems in front of tape

TAPE IS NOT DEAD – BUT IT IS FAST BECOMING A TECHNOLOGY TO AVOID – POOR TCO

# Implications on Storage Architectures

## It's all about the Servers

**Computer Architecture**
Von Neumann Architecture

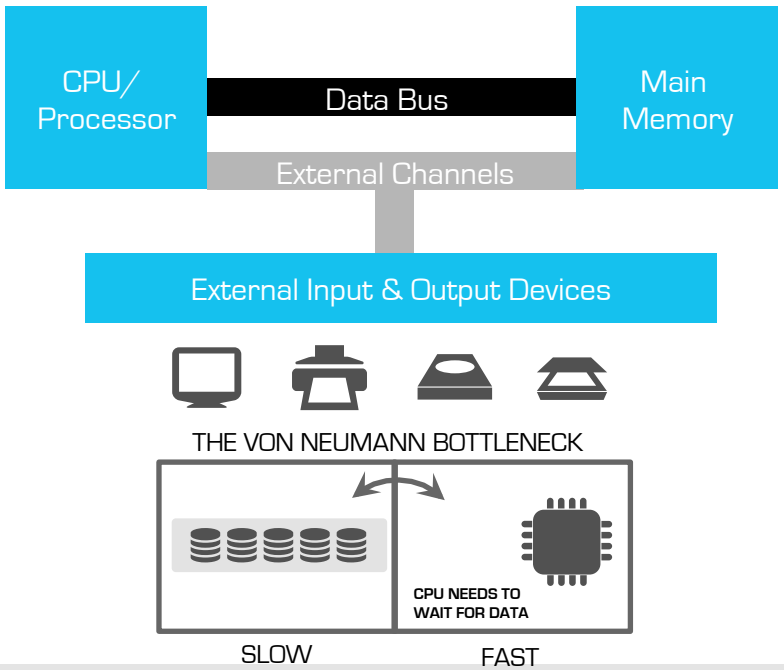


**In-Server Storage is not just Feasible but Optimal**

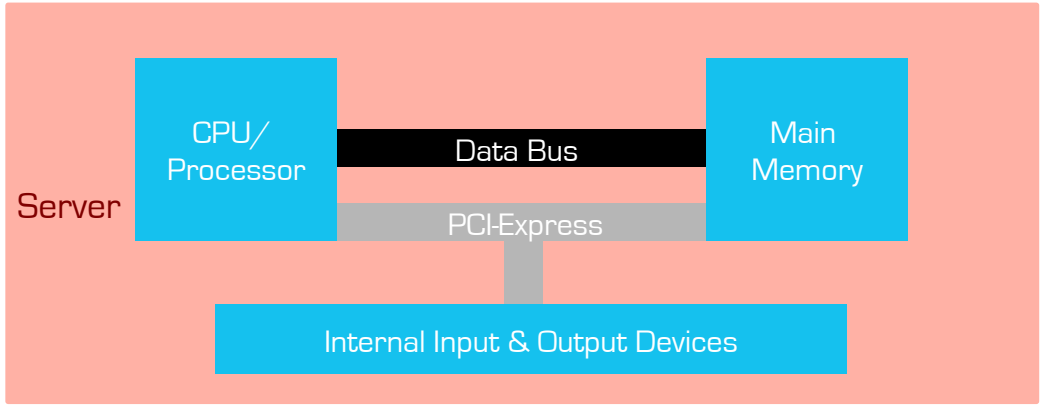


- Aggregate across servers – no more external Arrays
- Software controls the aggregation
- Built into all OSes today (e.g. Linux, Windows)
- Shared or non-shared – partitioning, clustering
- Very efficient scale-out (e.g. Ceph, GPFS, etc.)

**Combine Compute and Storage – Reduce Complexity and Management**
**You Need Compute Anyway – Might as well have it Close to Data**

# Implications for the Library of Congress - Example 1

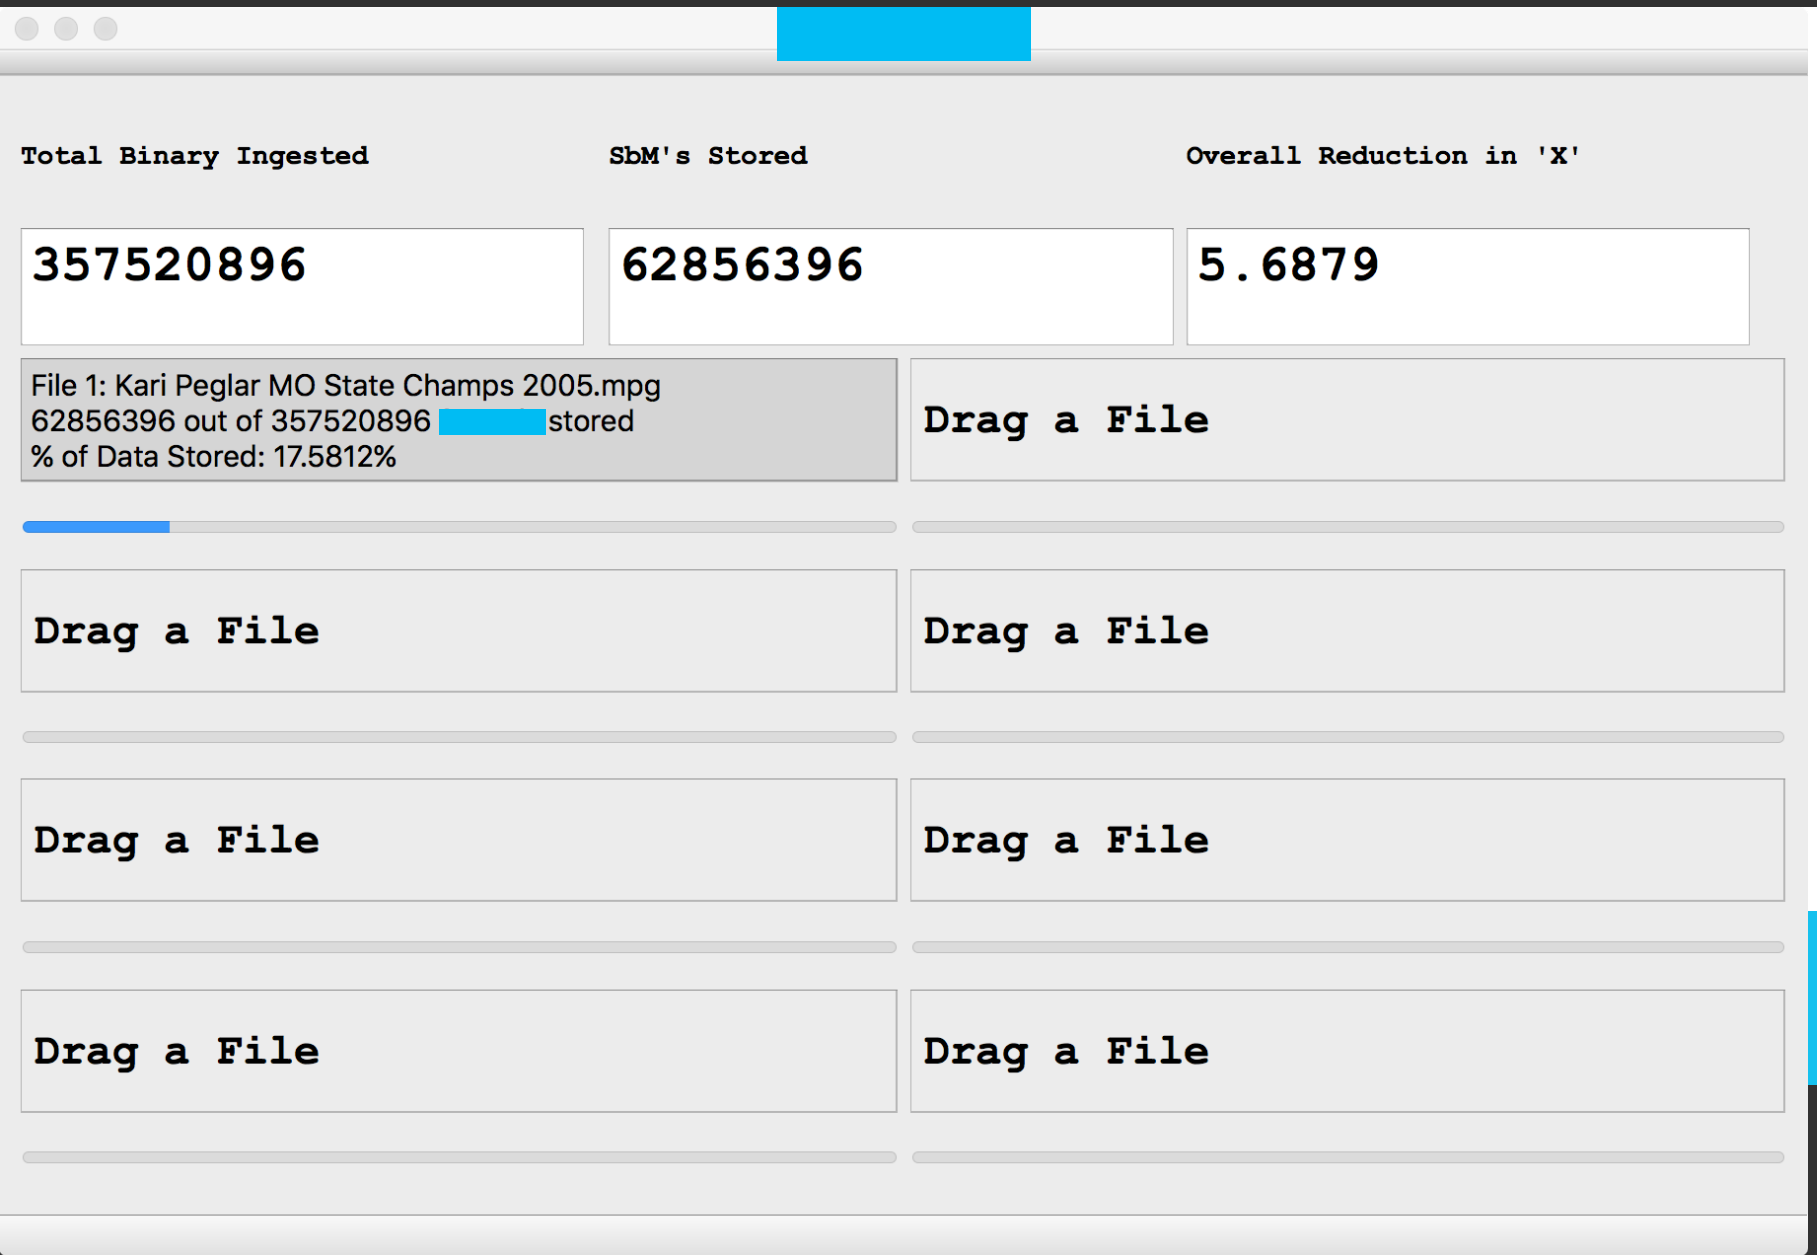


- 1.43 GB MPEG-4p2 file
- 357,520,896 4KB blocks
- Reduced down to
- 62,285,396 4KB blocks
- 251 MB on disk
- Amp factor 5.6879

# Implications for the Library of Congress - Example 2

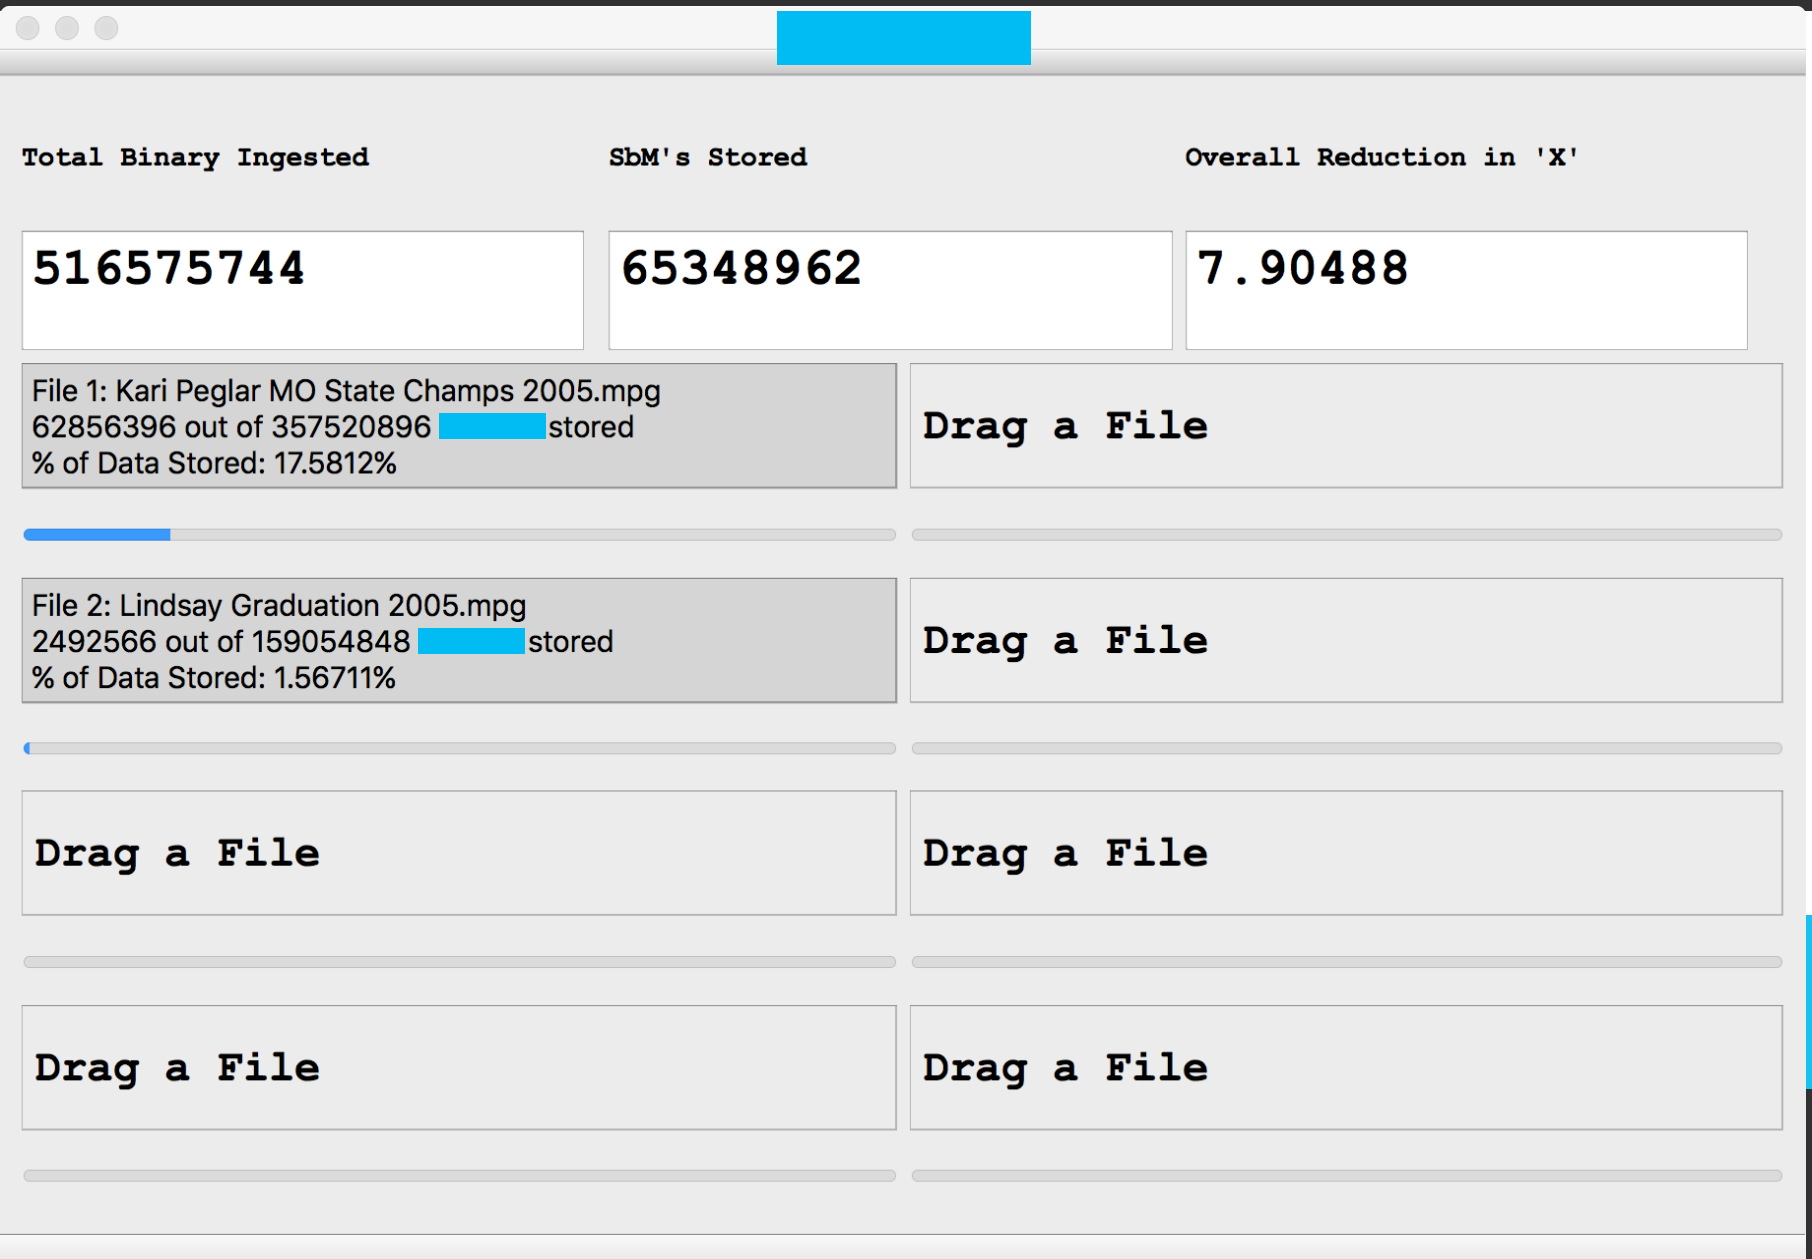


- Added a new 636 MB file
- 159,054,848 4KB blocks
- Total 516,575,744 blocks
- Reduced down to
- 63,348,962 4KB blocks
- 253 MB on disk
- Amp factor 7.90488

# Conclusion


**1 trillion**
connected objects an devices on the planet generating data in 2017

**2.5 billion**
gigabytes of data are generated every day

**44 ZB of data by 2020**
of which 80% is unstructured

**$ 266 B**
WW data spend by 2017

According to IDC