

Unintended COST'equences

Cost of not upgrading in a timely manner



Sean Leffler

HPC-SYS

03/13/2019




Managed by Triad National Security, LLC for the U.S. Department of Energy's NNSA

Supercomputing Storage @ LANL



100TB data sets per job!


Power user: 1PB/week x4 weeks
Only $\frac{1}{4}$ of system used




Campaign Storage (60PB)
Commodity drives
6months – 1year

Prior to Campaign +8PB/yr

2PB/yr to Archive



Tape archive



Current:

- 1 Oracle SL8500 Unclassified
- 1 Spectra TFinity Unclassified
- 6 Oracle SL8500 Classified
 - 1.5 PB disk tier

Future:


- 1 Spectra TFinity Unclassified
- 2 IBM TS4500 Classified

Classified Archive 70PB >>110PB
Unclassified Archive 10PB >> 18PB

50+ Years of data
Residence – forever
4% recall rate

Why are we conservative on tech

Gartner Hype Cycle

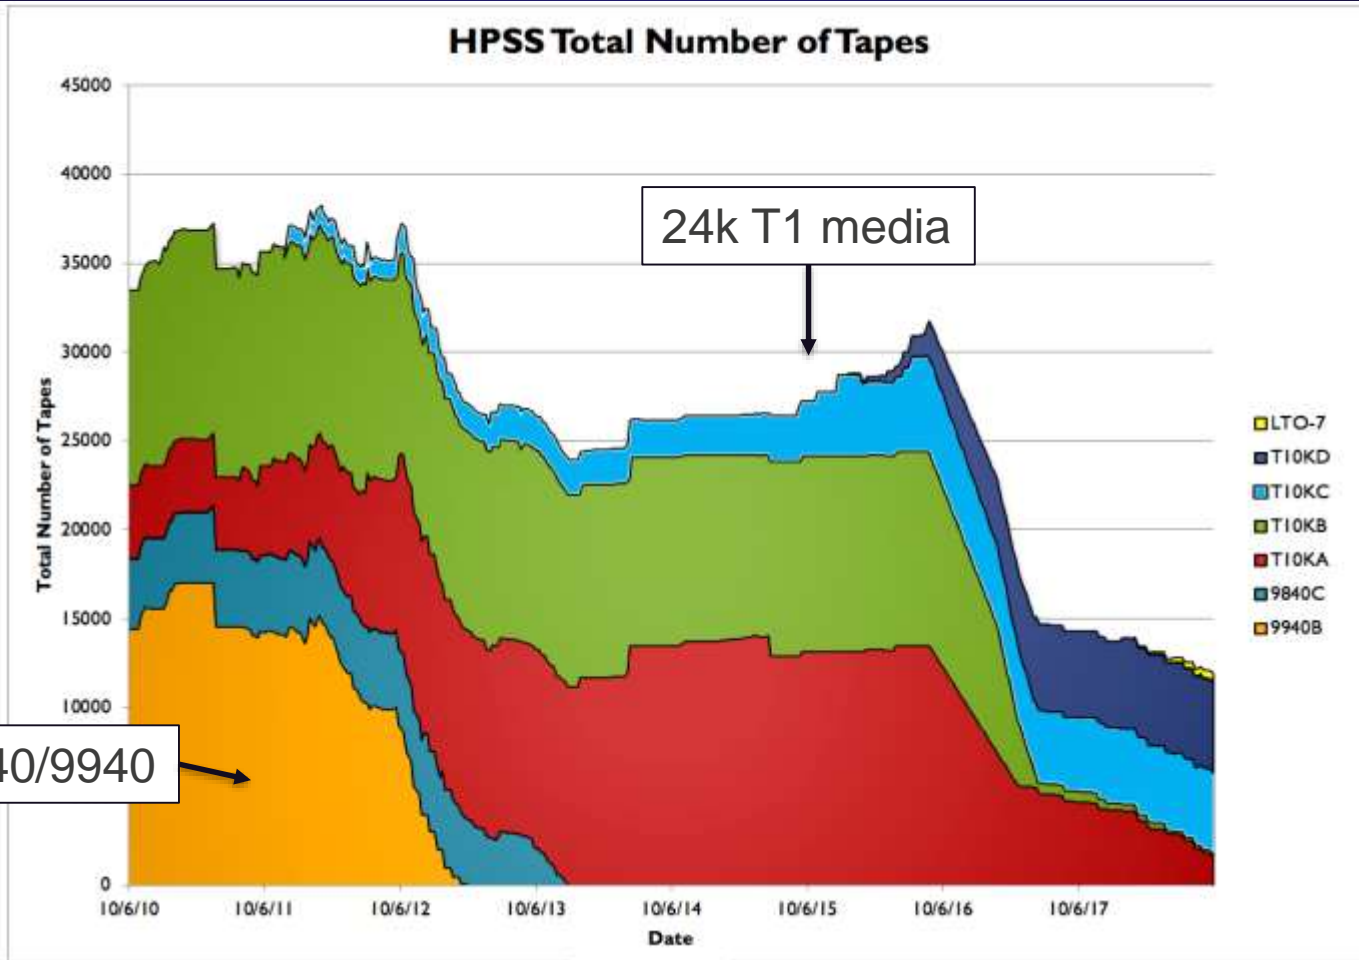


Why

- 50+ years of Unique Data
 - Hard/impossible to recreate
- Budgets
- Poor Marketing?
 - Hard sell, tech vs more tape

The result?

Results



The unintended costs

Labor!

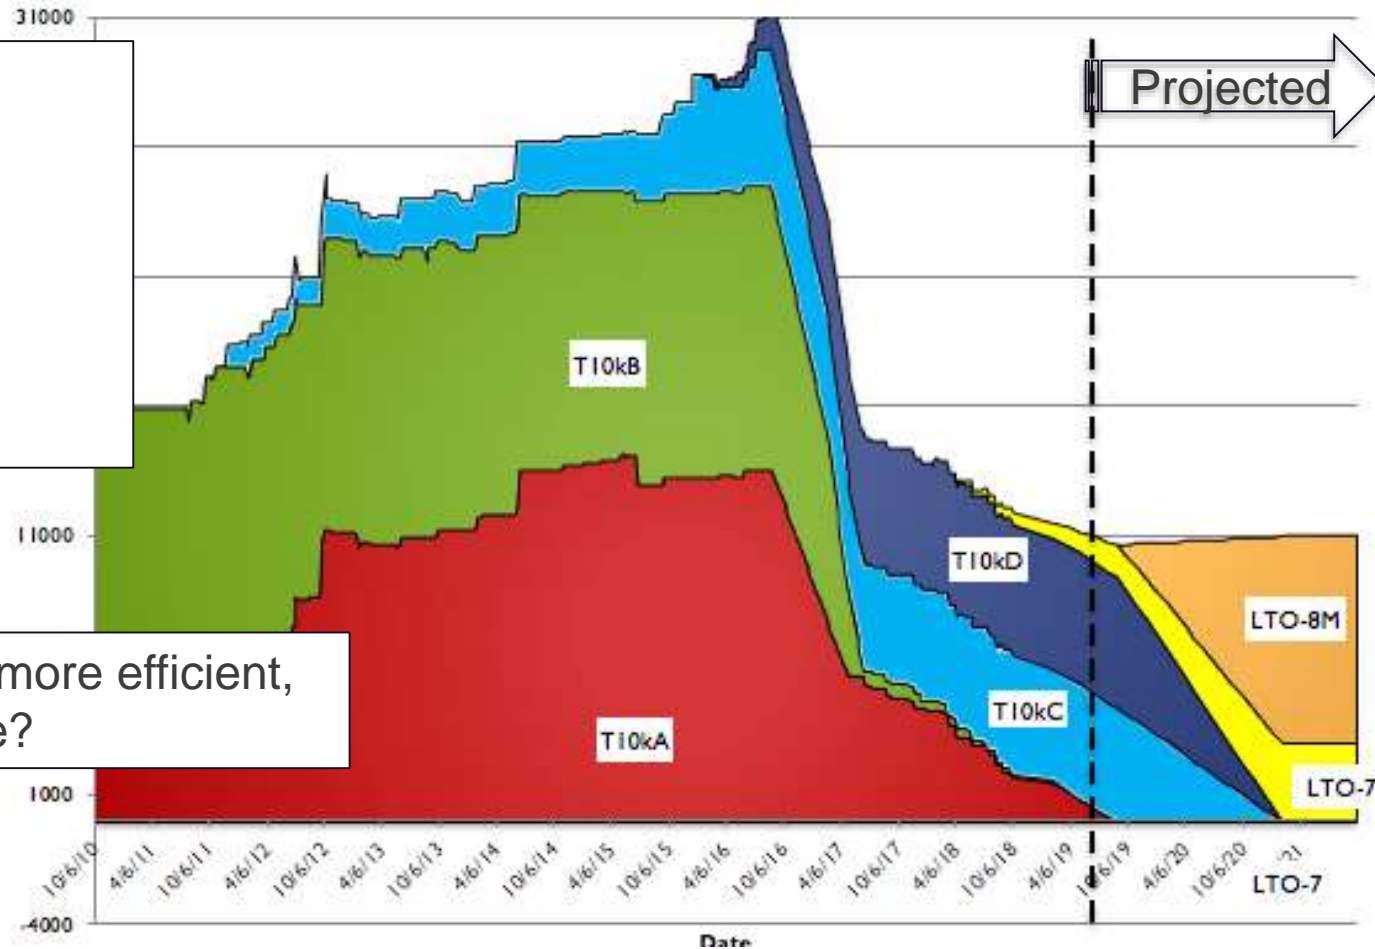
- 2+ years of effort to migrate old data
 - What is the cost to your company?
 - (Pay, benefits, etc.... The real cost to the company.)
- A lot of time spent maintaining the system
 - Damaged tapes require even more time to fix
 - Cause spurious errors.
 - Drives get disabled.
 - Congestion
 - On call getting paged
- Time not spent working on the future development

Future?

Classified capacity:
110PB 12,000 tapes
LTO-8M

Unclassified capacity:
18 PB 3,000 tapes
LTO7

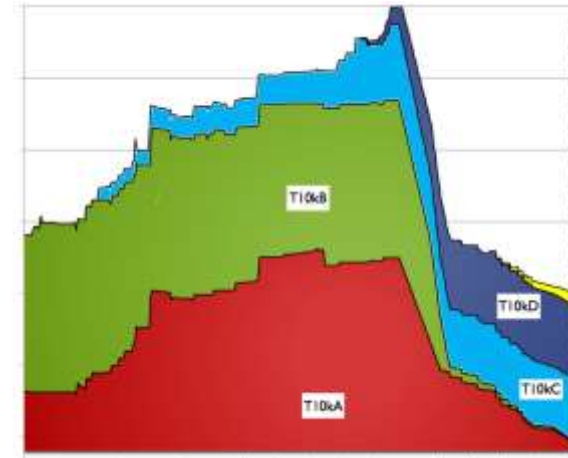
What affect will new, more efficient,
supercomputers have?



Projected

Solutions

- Don't let users save (everything).
 - Quotas (soft/hard)
- Purchase new technology
 - Gain capability! – media verification
 - Gain capacity! – denser media
 - Decrease footprint/power/cooling
- Don't wait (too long) to implement your new tech
 - Currently 6 months for testing/verification due to nature of data.



What if you can't buy new technology?

- Curate your data! Verify!
 - Move it around
 - Less of a concern with new libraries with media verification