
# NASA Nebula

## Cloud Storage for Open Science


I am not a scientist.


NEBULA


This is not your mother's storage array.


# Challenges

- Be all things to all people
- Be cheaper
- Be faster
- …and still be reliable (enough)

## Free Market


## Closed Market


# Approach

- Copy the private sector
- Ignore vendors
- Work from first principles
- Test, measure, and repeat
- Fail fast

# Find your cost leverage…



INFINIBAND™

# ..and match it.


# POSIX / POSIX-- / KVS

- Fast, temporary scratch disks
- “Pretty” big, “pretty” reliable NAS
- Very, very large KVS (by file size AND file count)


# Our conclusions

- Containerized Data Centers
- Open standards: HTTP, WebDAV, Ethernet
- Monitoring-in-depth
- Do battle with every SPoF
- Recompute

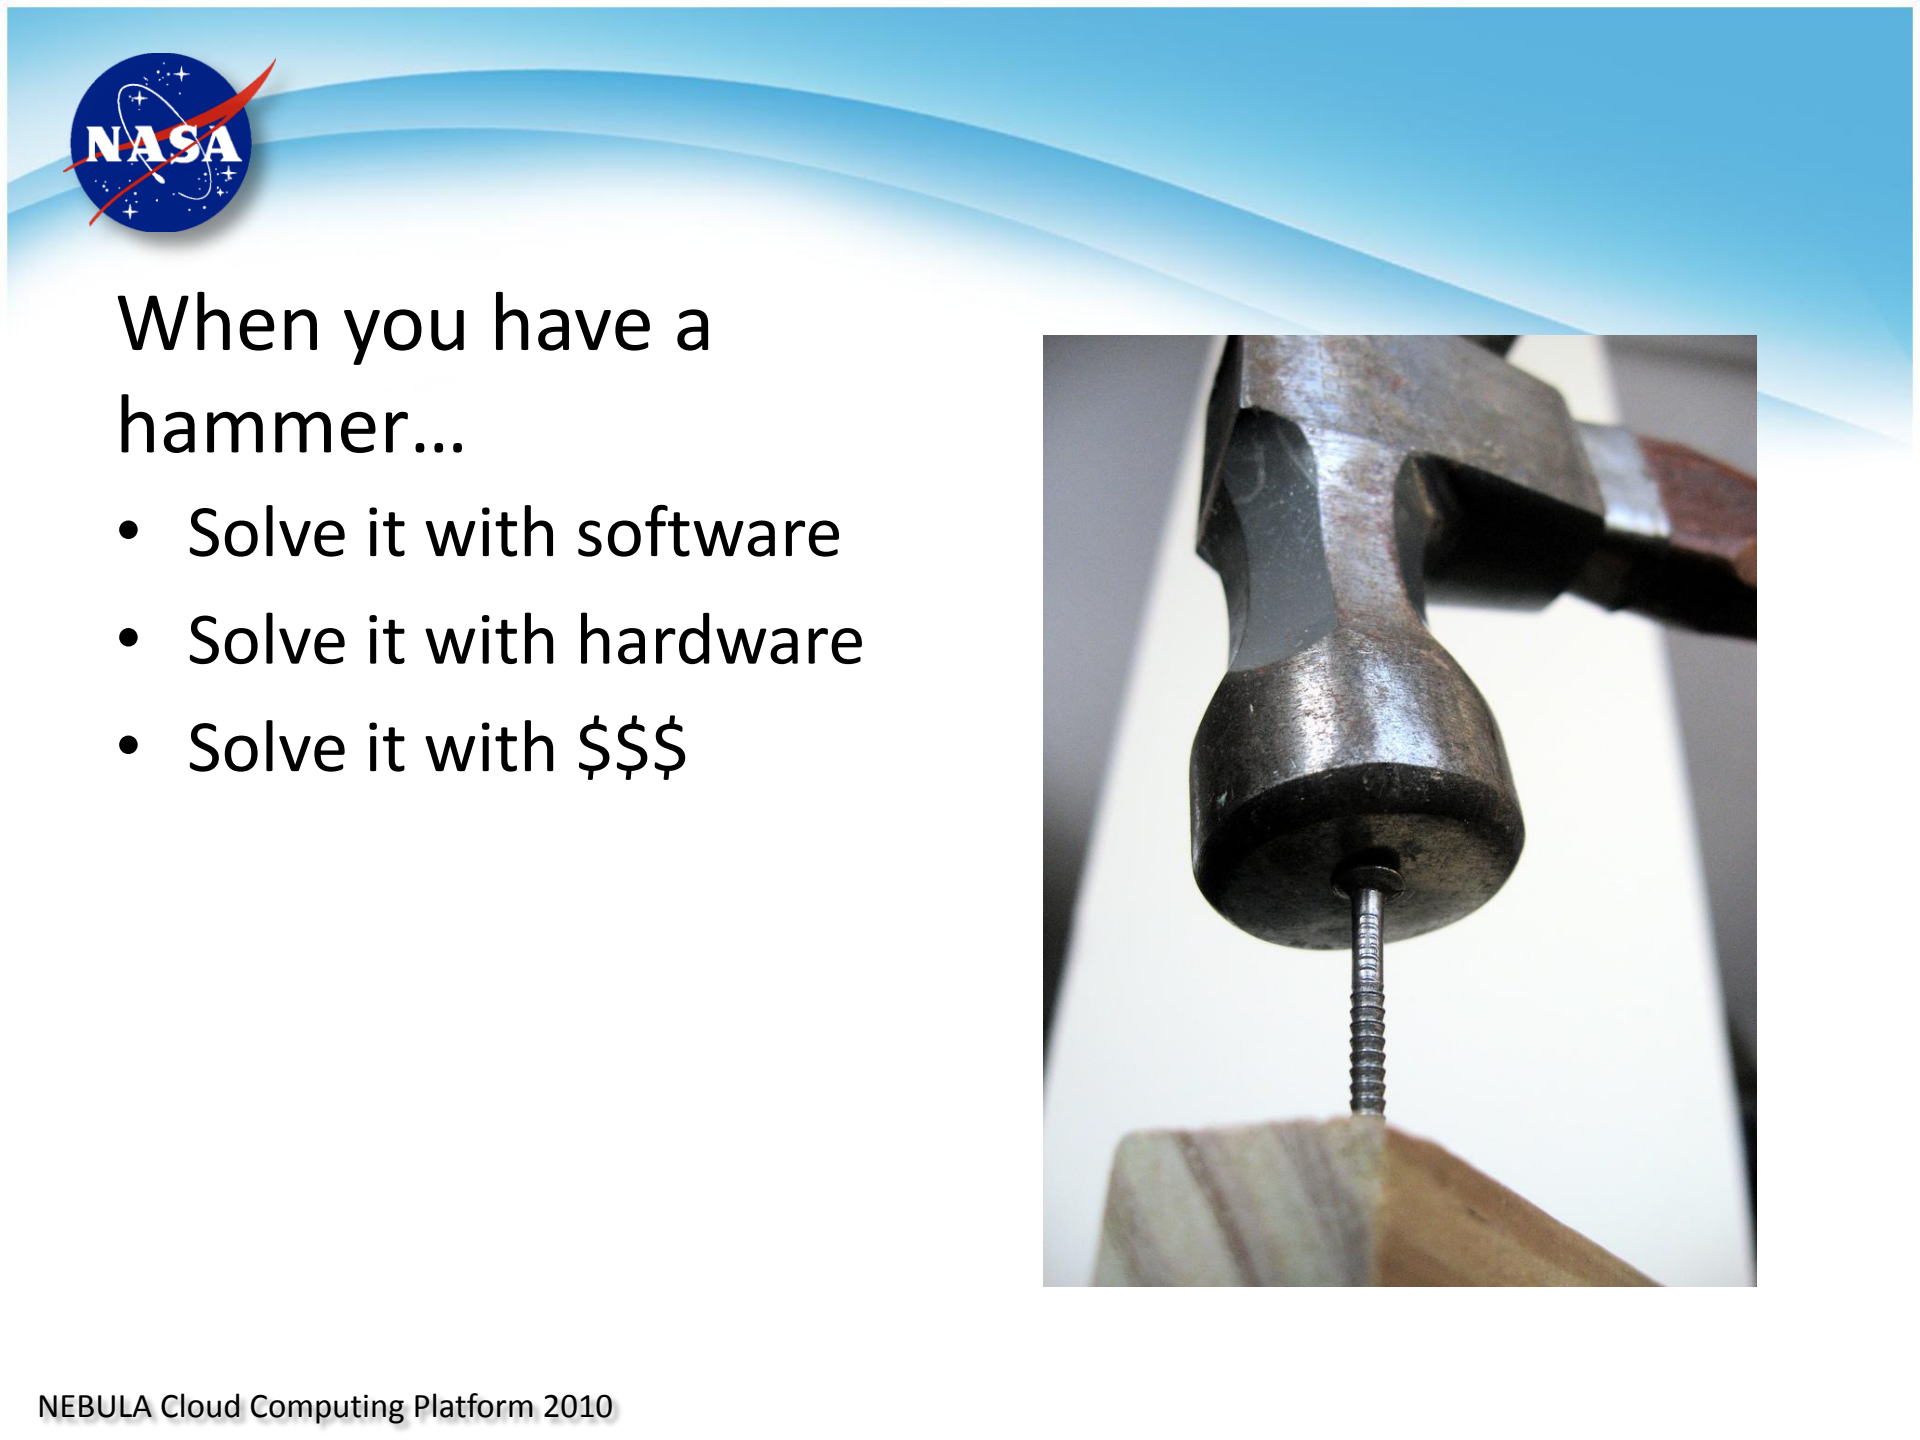


# When you have a hammer…

- Solve it with software
- Solve it with hardware
- Solve it with $$$


# Custom solutions (ongoing)

- Evented, parallel web service interface
- Multicast file replication (IPv4 and IPv6)
- MapReduce-supporting Object Store
- Alternate filesystems
- FusionIO

# The Gauntlet


## Your DATA
- Raw Data
- Metadata

## Storage Interface
- POSIX
- Web Services

## NAS Protocol
- AoE (L2)
- iSCSI / NFS (L3)

## File System
- Distributed (Stripe Size)
- Local FS (Block Size)
- Partition Table (maybe)

## RAID Volume
- Block size
- Drive Count

## RAID Controller
- BBU
- Cache Size

## Raw Disk
- Interface (SAS / SATA / SSD)
- Reliability


# Stop passing the buck

- Stated risk: 10^16
- Actual risk: 1 in 72


# References

- http://lauraparkin.wordpress.com/2009/07/11/fail-early-fail-fast/
- http://nebula.nasa.gov