

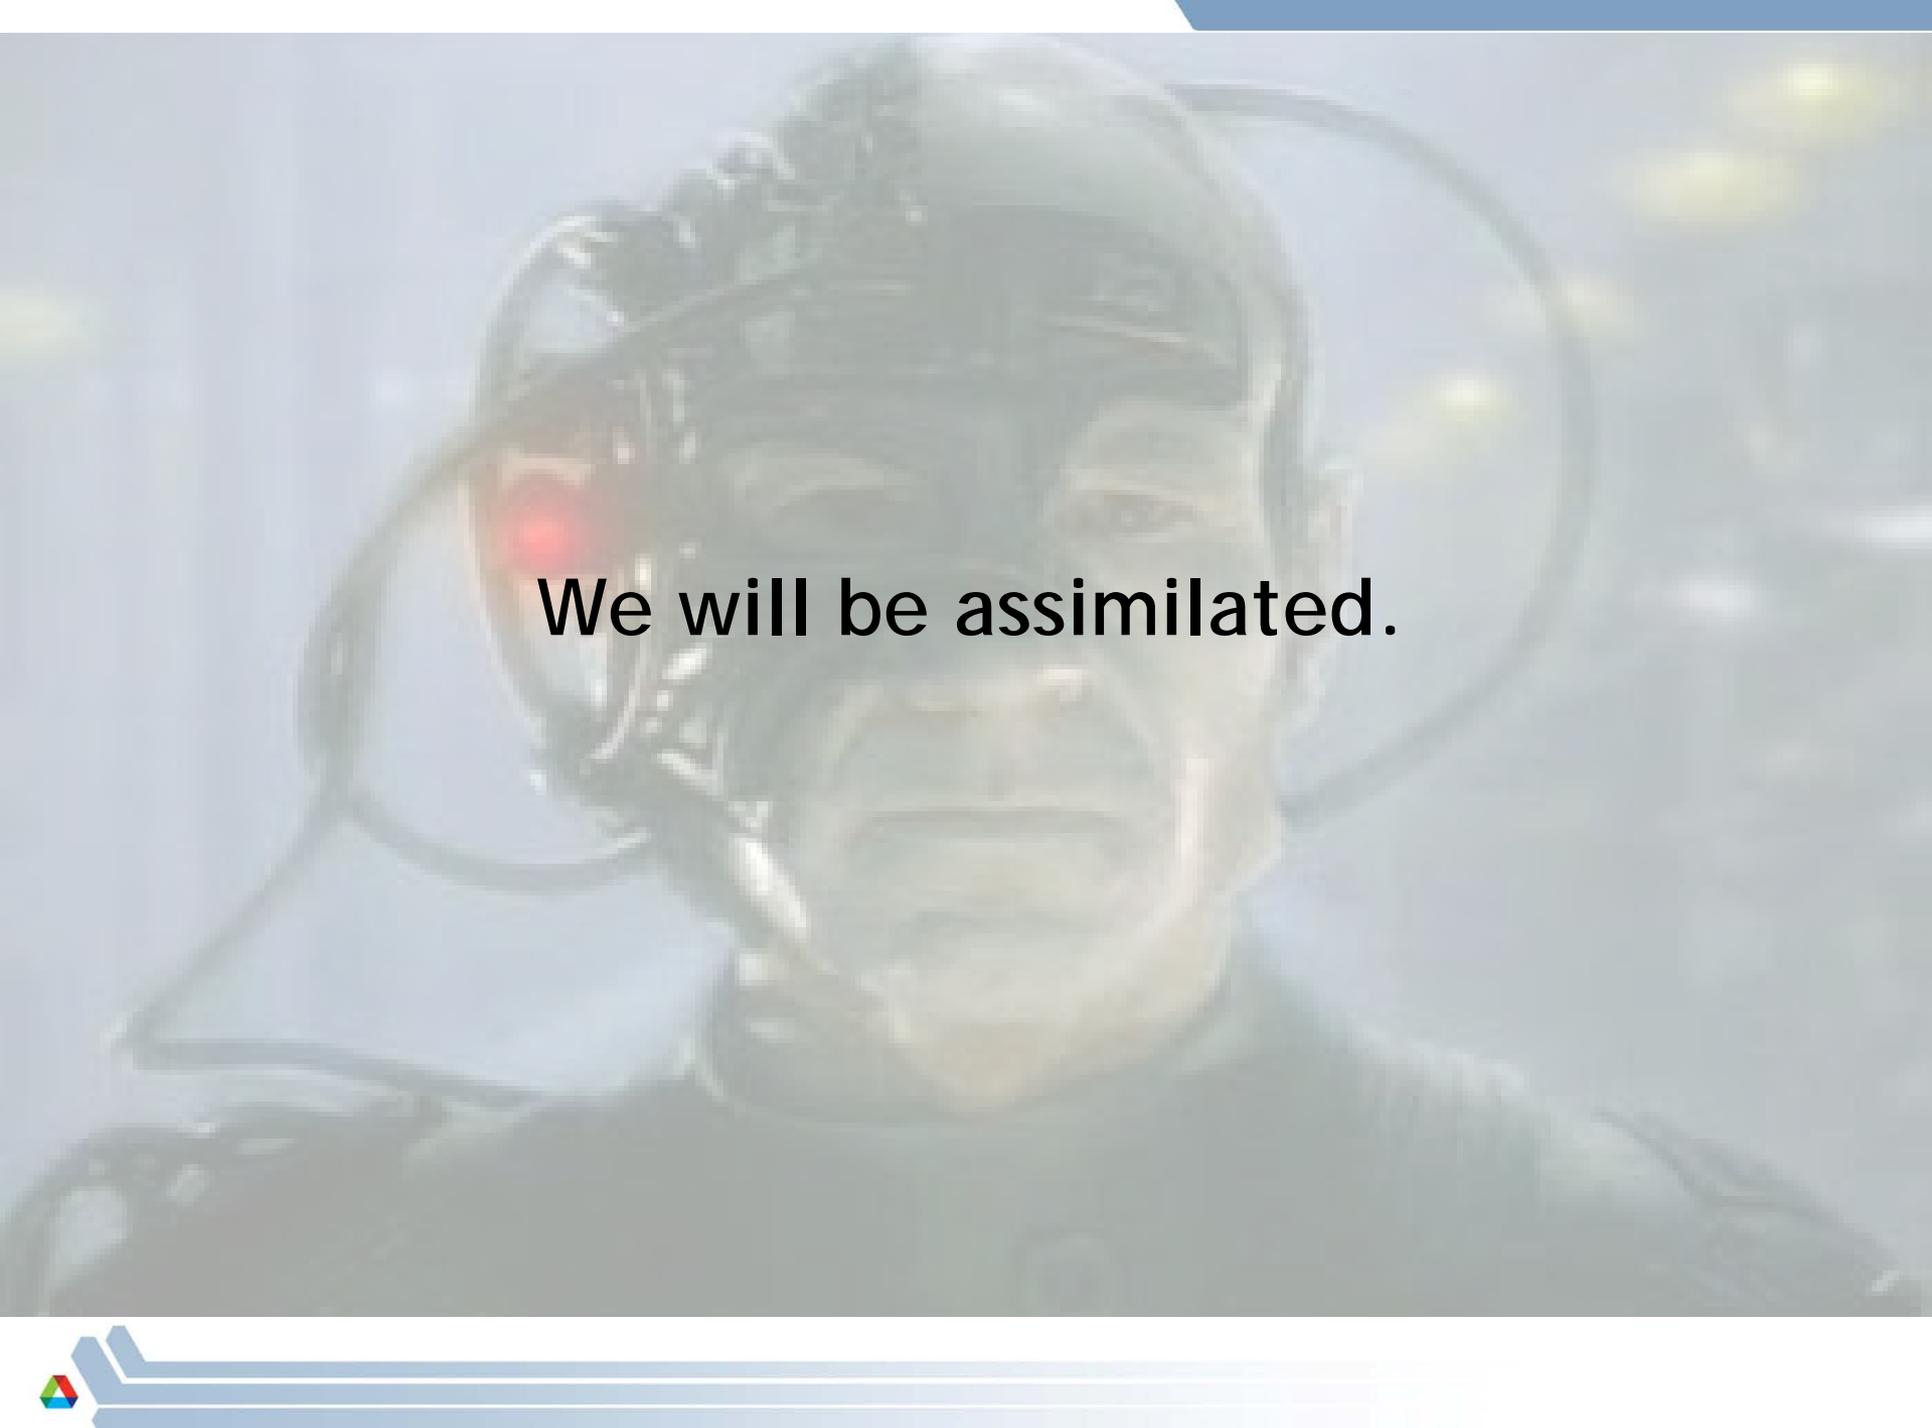
The Future

Rob Ross

Mathematics and Computer Science Division

Argonne National Laboratory

rross@mcs.anl.gov

A close-up, slightly desaturated image of a man wearing a futuristic, dark helmet with a prominent red light on the left side. The man's face is partially obscured by the helmet's structure. The background is a blurred, light blue-grey color with some out-of-focus lights. The text "We will be assimilated." is centered over the man's face in a bold, black, sans-serif font.

We will be assimilated.



Hadoop will supplant all our parallel file systems.

- Software is free and open source
- Big community of users, developers, and support staff
- Already exists
- Operates on commodity, shared-nothing storage
- Fault tolerant
- Low coordination cost during concurrent I/O (using PLFS or similar)
- Exposes locality of data, supports locality-aware scheduling



We should help improve Hadoop and adapt our tools to welcome our new data overlords.

- **Distributed metadata** – help support multiple NameNodes
- **Networking** – support HPC networks
- **Concurrent writes** – control data placement at write time to ensure high bandwidth checkpointing (separate files)
- **Failure detection** – heartbeat to NameNode must be replaced
- **Intermediate storage** – integrate burst buffers into write path
- **Data model** – implement middleware to support scientific data models on top (e.g., BigHDF5, BigAMR, or BigClimate)
- **Analysis tools** – port analysis tools to operate through new data model middleware





See you next year at HEC-DISC-FSIO 2011!

