

Hi. I'm Wayne Sawdon from Almaden Research. My name is Jim Gerry from IBM Houston.

We're here to show the Billion File Demo. We stored a billion files on the GPFS File system. Not only can we store a billion files, but were going to process them with the ILM (Information LifeCycle Management) policy and find out what's new and what's changed. We're going to run a backup scan with a billion files and find everything that's changed in about 5 hours. In research labs we can do it faster but on the floor here we can do it in 5 hours. Once I find these files that changed I'm going to hand them to HPSS.

What HPSS does is we HSM manage this GPFS file system. That allows you to store petabytes of information on to a file system with merely a handful of terabytes of disk drives. In addition to that we utilize the Island policy scan to back up the GPFS file system. We can backup the file system, collect the I-nodes and the extended attributes. We also store that information into HPSS and that allow us to restore the GPFS in the case that it fails due to some kind of catastrophic accident.

Together we are showing file system support for a billion files.