GPFS Updates: what's new in v4.1

Vladimir Sapunenko (CNAF)

Notes from GPFS user Group meeting

April 29th, London, UK

New major release: 4.1

- Available from the second week of May
- Several performance improvements
- new NSD format
  - NSD created in 4.1 is not compatible with previous versions of GPFS
- Better timeout management
  - automated deadlock detection, data collection and breakup
- Improvements in Monitoring
  - Plug-in for “dstat”
- Encryption (new)
- Integration with OpenStack
  - OpenStack “Havana” and “IceHouse” include GPFS Cinder driver

New features in AFM

(Active Files Management)

- Support for parallel I/O
  - Multiple gateways can transfer a single (big) file in parallel
- Prefetch
  - handling gateway node failure durin prefetch
- Native GPFS protocol support
  - In addition to NFS, GPFS protocol can be used for AFM communication and transfers
  - Integration of GPFS features and attributes
- AFM based migrations from any legacy NFS systems
  - Minimize downtime while migrating data to a new HW
  - Update of GPFS file system format (ex. Enabling 4MB block size)

Other new features

- Clustered NFS (**cNFS**)
  - NFSv4 and IPv6 support
- **Quota** management
  - enable and disable on-line
  - quota files are metadata now (no user.quota, group.quota files in FS)
- Local read-only cache (**LROC**)
  - Overflow pagepool to local storage (SSD)
  - more RAM for application
- **FPO** - file placement optimizer (Hadoop-like)
  - data locality when restriping
  - AIO performance
  - mmchpool to change FPO attributes
  - Mixed storage pools in one file system

Other new features (cont.)

- **CCR**: Cluster Configuration Repository
  - quorum based instead of server based
- **RDMA** over Ethernet (will be available as an update later)
- **Fileset** snapshot restore
- **4k** Disk sector support
- **Cyrwin** replaces **SUA** (Subsystem for UNIX-based Applications) on Windows nodes
- Message logging with syslog
- On-line migration of Extended Attributes (**EA)**
  - Mmmigratefs can be run on mounted FS
- **Rapid repair** of large replicated files when restarting disk
  - Repair only block which were changed while disk was down
- User-defined node classes

Encryption

- Native encryption support in GPFS
- Encryption of data at rest
- HIPAA, Sarbanes-Oxley, EU, and various national data privacy laws compliance
- Multiple keys
- Secure deletion

New licensing model

- Three levels of functions
  - Express Edition
    - basic GPFS features (no HSM, cNFS, AFM)
  - Standard
    - = 3.5
  - Advanced
    - Native encryption for secure storage and deletion
  - All nodes in a cluster **must** be at the **same level**

Questions?