EMC Lustre Contributions

It's all about speed.

Tao Peng
Xuezhao Liu
as presented by John Bent
Fast Data Group
Office of the CTO

EMC²

EMC Lustre activities

- Support Lustre bug fixes (LU-1126, LU-1322, etc.)
- Lustre client mainlining work:
  - Clean-up the Linux Lustre client for inclusion in Linux kernel
  - Port Lustre client to latest kernel for upstream inclusion
  - 37 patches accepted, 36 patches under review, more to come in.

EMC²

EMC Lustre activities

- Write a new IETF draft for pNFS Lustre specific layout – discussed at latest IETF meeting: Sorin Faibish+Peter Honeyman leaders
- Introduce a new Lustre specific benchmark component in the SPECsfs 2013 – approved at SPECsfs meeting: Sorin Faibish leader

EMC Lustre activities: VNX HPC Series

- High IOPS/ Throughput
- Small Form Factor
- High Density
- Best Price/ Performance
- Enterprise Reliability, Availability and World class Service

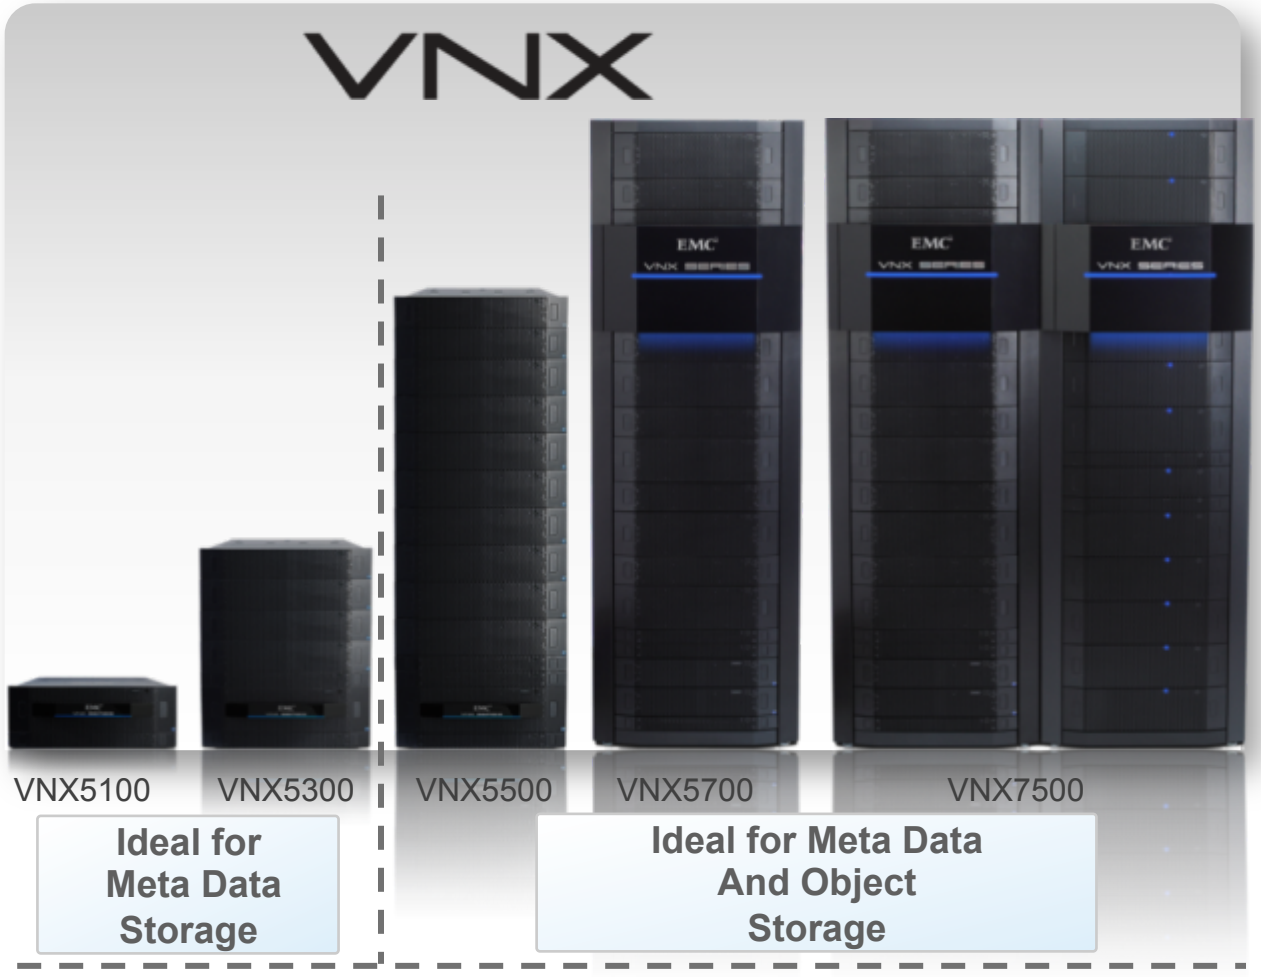

EMC²

EMC Lustre activities: VNX HPC Series

- Base Configuration is a Single Rack offering
  - 720 TB Capacity, 8 GB/s Performance
  - Pre-racked and configured VNX5100 and VNX7500
  - Servers for Management and File System
- Single Point of Management via Management Console from Terascala
- Application Ready - Pre-configured and tuned Lustre Parallel File System
- Infini-band (QDR) interface to computational node

EMC²

Lustre Client Mainlining Work

- LSF/LUG Discussions
- Current Status
- Future Work
- User Impacts

EMC²

Mainlining Lustre Client: LSF/LUG Discussions

- General agreements
  - Lustre client in the kernel is helpful for users and Lustre community
  - Must write code in the kernel way
  - Lustre is complex. It takes time to cleanup: any contribution from Lustre community is well received

Lustre Client Merge Requirements (LSF)

- Old kernel support is not allowed in the kernel
- Different platform support is disliked by most attendees. (Except for Al Viro who wanted to see code first but still no response)
- Server code can be marked out with HAVE_SERVER_SUPPORT

EMC²

Mainlining Lustre Client: Suggestions (LSF)

- Put in staging tree to clean up and get merged
  - Need to be compliable
  - Need to be properly signed-off
- Steps:
  - Huge patches to merge in staging tree
  - Put Lustre code in fs/ directory and depend on staging
  - Small and incremental patches to cleanup
  - New feature patches are allowed
  - Get each piece of code properly signed-off by key reviewers then the code can be truly merged

EMC²

Mainlining Lustre Client: Current Status

- Clean-up the Linux Lustre client for inclusion in Linux kernel
  - coding style changes (on-going)
  - cleanup old kernel configure checks (done)
  - cleanup ptlrpc/libcfs code (almost done)
  - cleanup compat code (on-going)
  - split server/client code (done)
- Port Lustre client to latest kernel for upstream inclusion
  - 3.0 support done together with Cray
  - 3.4 kernel support patches under review
  - 3.6 (latest) kernel support under development

EMC²

Mainlining Lustre Client: Future Work

- Continue Clean-up
- Continue with newer kernel support
- At a proper point
  - Use proper scripts to cleanup leftover libcfs and old compat code
  - Change Lustre source code layout (out of Whamcloud tree) and put it inside Linux kernel to build
  - Submit to upstream kernel for review

EMC²

Mainlining Lustre Client: User Benefits

- When Lustre client merged in Linux kernel:
  - Liunx kernel has much larger (user/developer) community
  - Bring more users and developers to Lustre
  - More usage/deployment results in more stable Lustre client
  - More development to win more fancy features
  - Much easier (and quicker) to support newer kernels
  - Take advantage of new kernel features more quickly

EMC²

Mainlining Lustre Client: User Benefits

- When Lustre client enters Linux vendor distributions:
  - Avoid tainted kernel with Lustre in vendor distros (tainted kernel means limited vendor support)
  - Use Lustre client with official support from your Linux distribution vendors

john dot bent at emc com

wc-discuss mailing list