



File System Technologies in a Linux Cluster Environment



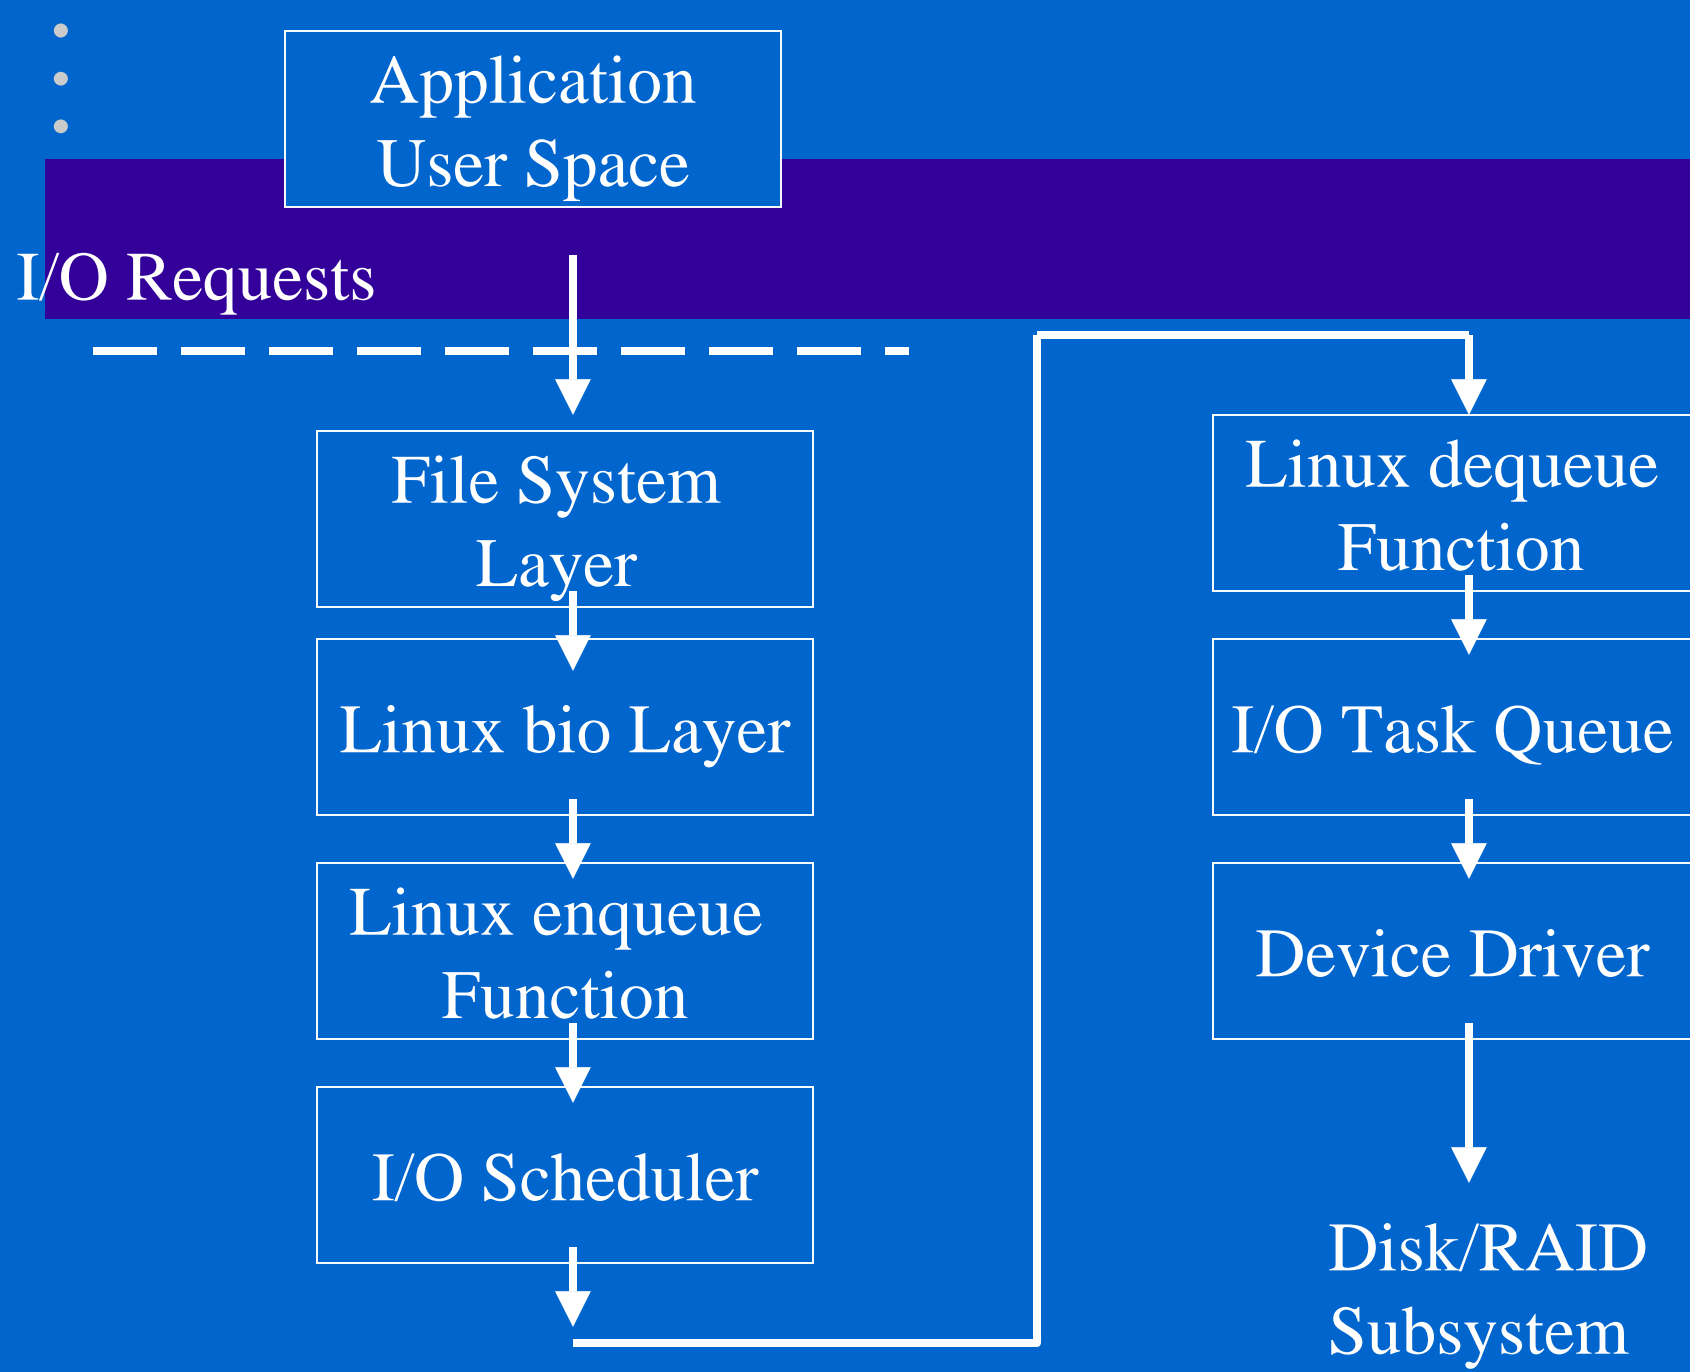
Dominique A. Heger
DHTechnologies (DHT)
www.dhtusa.com



-
-
-

Agenda

- UNIX/Linux OS Abstraction (PNS, DNS)
- Linux 2.6 I/O Subsystem
- Cluster File System Terminology
- IBM's GPFS
- Red Hat's GFS/GFS2
- Q&A



•
•
•

Cluster File System Terminology


- Distributed file system
- Global file system
- SAN file system
- Symmetric file system
- Asymmetric file system
- Parallel file system

-
-
-


IBM's GPFS

- Powers some of the largest HPC clusters
- 2PB file system in production
- 122GB/sec (read & write) throughput
- Striping data across n servers and m disks
- Client side caching, configurable block size
- Large directory support (EH)
- Data, file, and file system metadata locking


IP Network




GPFS Nodes




Disk Subsystem



GPFS
IO Nodes



Disk
Subsystem

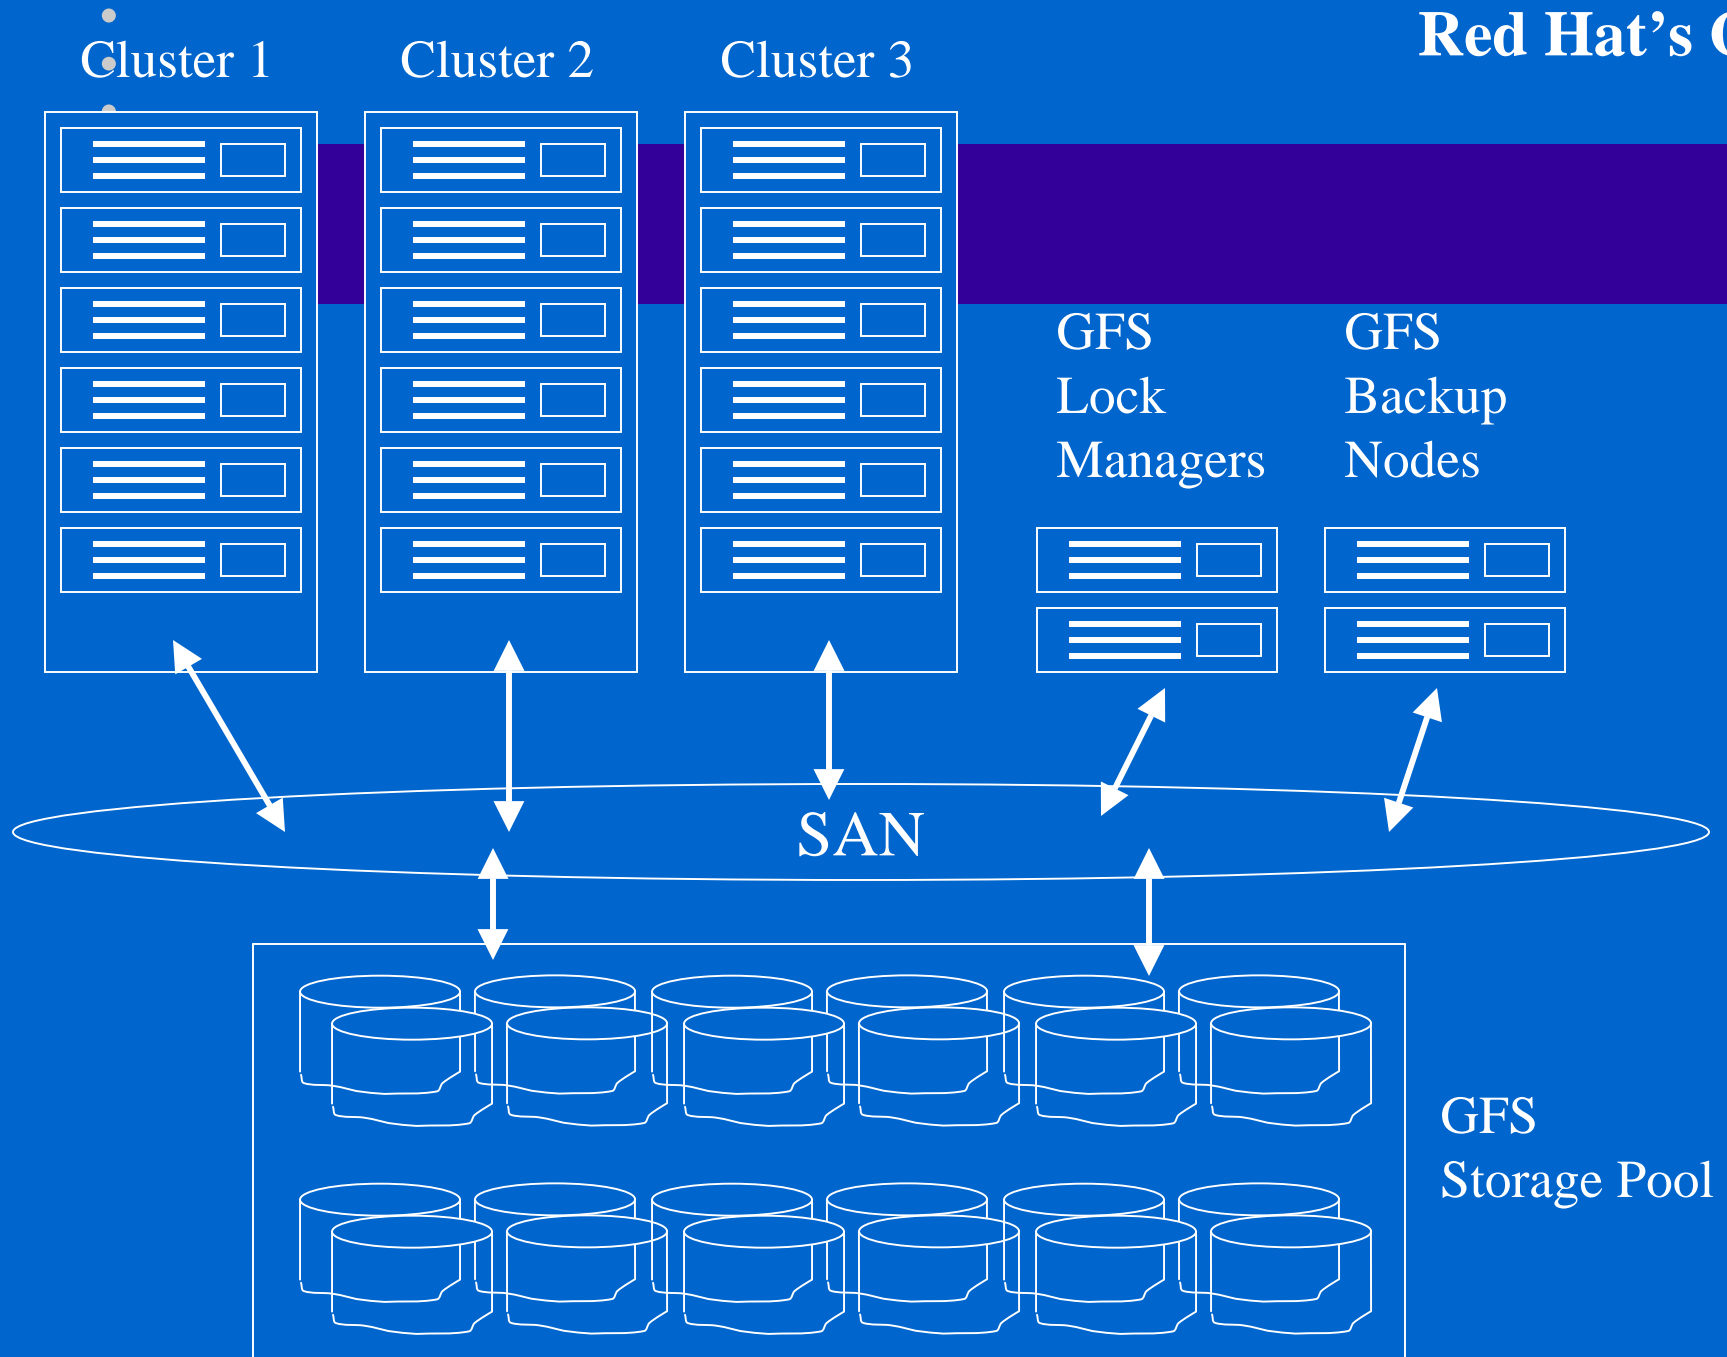


•
•
•

Red Hat's GFS/GFS2

- 64-bit parallel file system
- Journaling file system (AWB)
- Transaction manager & Log manager
- NSPVD (GNS, striping), CLVM (mirroring)
- dlock(), dmep(), SLM, RLM, RH-DLM
- Extendible hashing for directories
- Direct IO, ACL, metadata file system (GFS2)


Red Hat's GFS



-
-
-



Q&A



-
-
-
-
-
-
-
-