SNIA DSI '14
DATA STORAGE INNOVATION CONFERENCE
APRIL 22-24, SANTA CLARA, CA

Project Fermi - A Highly Available NAS Gateway Built from Open Source Software

Dan Pollack

AOL


Enrico Fermi

SNIA™

Why NAS gateways?

- Decouple storage from NAS server
- Widely deployed SAN infrastructure
- Storage arrays cost less than captured storage in NAS servers (@AOL)
- Multiple tiers of storage in NAS servers
- Heterogeneous storage in NAS servers

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What is Fermi?

- Commodity: Industry Standard Servers (ISS) and Open Source Software (OSS)
- Redundancy: Clustered heads for redundancy
- FC backend: Performance and reliability

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

Why Fermi?

- Cost Savings
  - Cost of commercial NAS servers - $$$$
    - Expensive hardware
    - Software licenses
  - Cost of Fermi - $
    - ISS hardware
    - No licenses
- More Flexibility (Any storage OK)
- More rapid hardware refresh

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What hardware?

- 64-bit x86 CPU server
- Multiple 1Gbps/10Gbps Ethernet links
- Multiple 8Gbps/16Gbps FC links

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What software?

- Open Solaris (open Indiana)
- Pacemaker (HA cluster resource manager)
- Heartbeat (HA cluster infrastructure manager)
- ZFS
- Solaris CLI

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What did we add?

- FC disk cluster SCSI reservation management
- ZFS zpool management
- Virtual IP address management
- Minor solaris specific bug fixes in pacemaker

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What is complete?

- HA Failover
- NFS services
- iSCSI services
- SNMP metrics

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What needs doing?

- CIFS services
- Active Directory integration
- Local snapshots
- Remote snapshots (ZFS send/receive)
- Provisioning API

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What does it look like?

- Three servers in a stack
- Multipath FC
- Multiple Ethernet interfaces
  - Multi-home
  - LACP/MLAG
- HA services for network and storage
  - Virtual IP for network failover
  - Shared storage for failover

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

What does it look like?

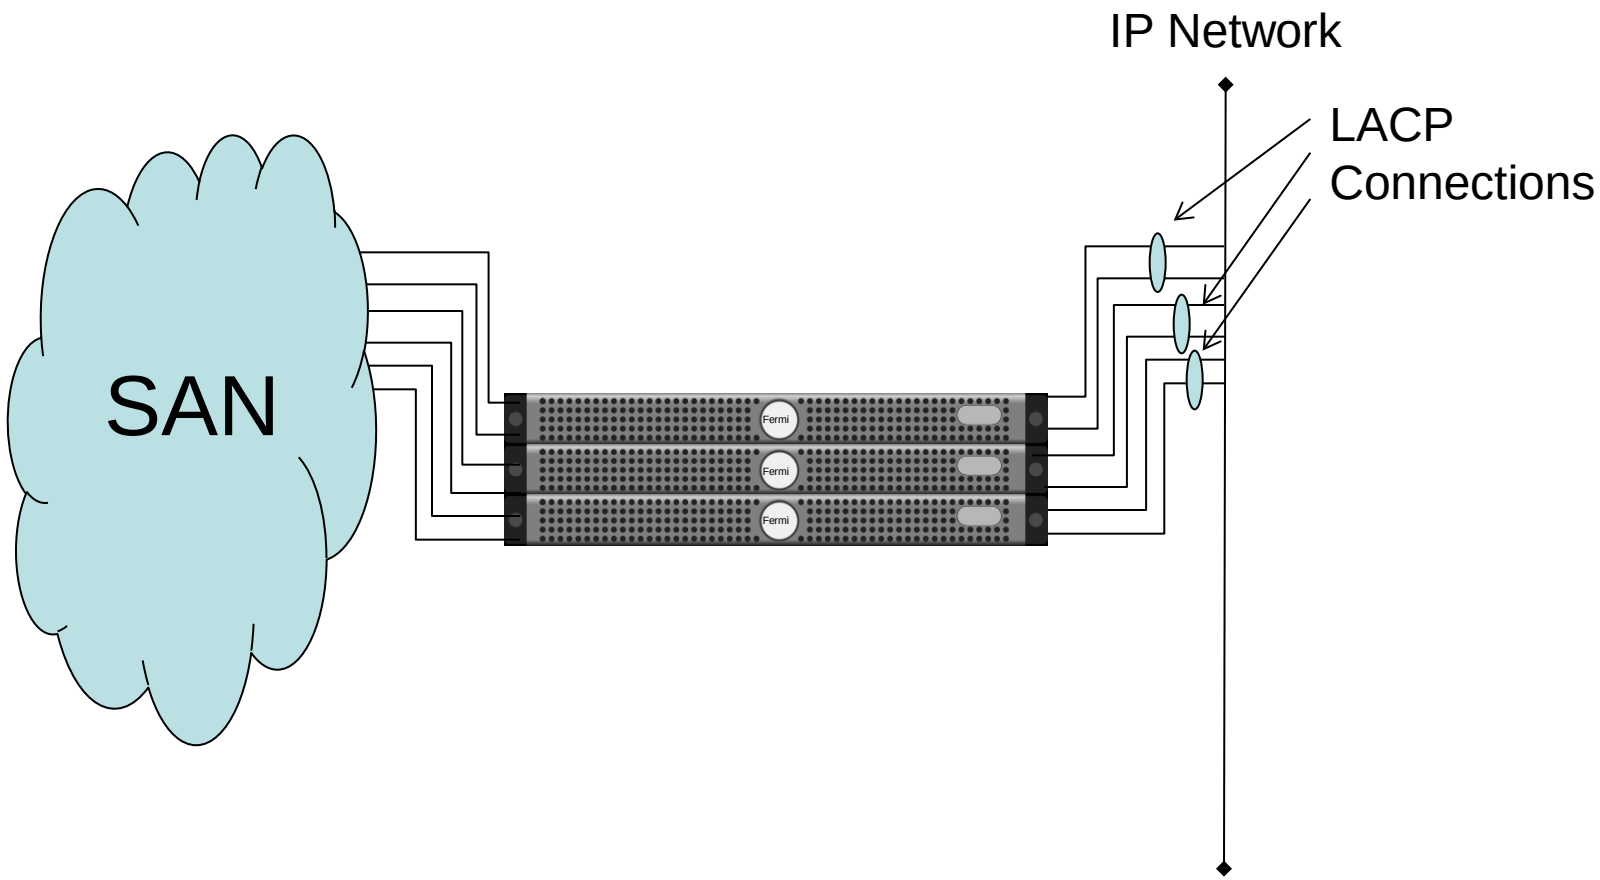

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

How does is perform?

- Filebench results

test	IOPS average	IOPS median	IOPS min	IOPS max	Bandwidth average MB/s	Bandwidth median MB/s	Bandwidth min MB/s	Bandwidth max MB/s
fileserver	7098	7166	4695	7907	169	171	112	188
webserver	8090	8063	7719	9053	40	40	38	45
randomwrite	8559	8755	5393	9466	66	68	42	73
openfiles	8518	8471	7916	9397	0	0	0	0
randomrw	9846	10162	3897	12494	76	79	30	97
mongo	6988	7004	3502	7005	22	23	11	23
randomread	8920	11259	70	12066	69	87	0	94
varmail	2931	3296	358	5627	10	11	1	20

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.

Questions?

2014 Data Storage Innovation Conference. © AOL. All Rights Reserved.