

open



USE



IMPROV
E



EVANGELIZ
E

Open Storage

Dominic Kay
Sun Microsystems

開
放
的
열린
مفتوح
libre
मुक्त
ಮುಕ್ತ
livre
libero
ముక్త
开放的
açık
open
nyílt
•••••
πικρ
オープン
livre
ανοικτό
offen
otevřený
öppen
открытый
வெளிப்படை



Agenda

- Why
- What
- Examples
- The Toolbox
- Call to Action
- Q&A



Today's Storage Architecture

- SAN
- DAS
- NAS



Changing Approach To Storage

- General purpose storage building blocks
 - Scalable, secure, open operating systems
 - Data friendly industry standard server platforms
 - Commodity disk arrays
- Open standards for storage and data management
 - Open Source
+ Industry Standards
= Open Storage
 - Open source in order to continue to propagate the standards, and ensure high quality, interoperable, flexible cutting edge solutions.

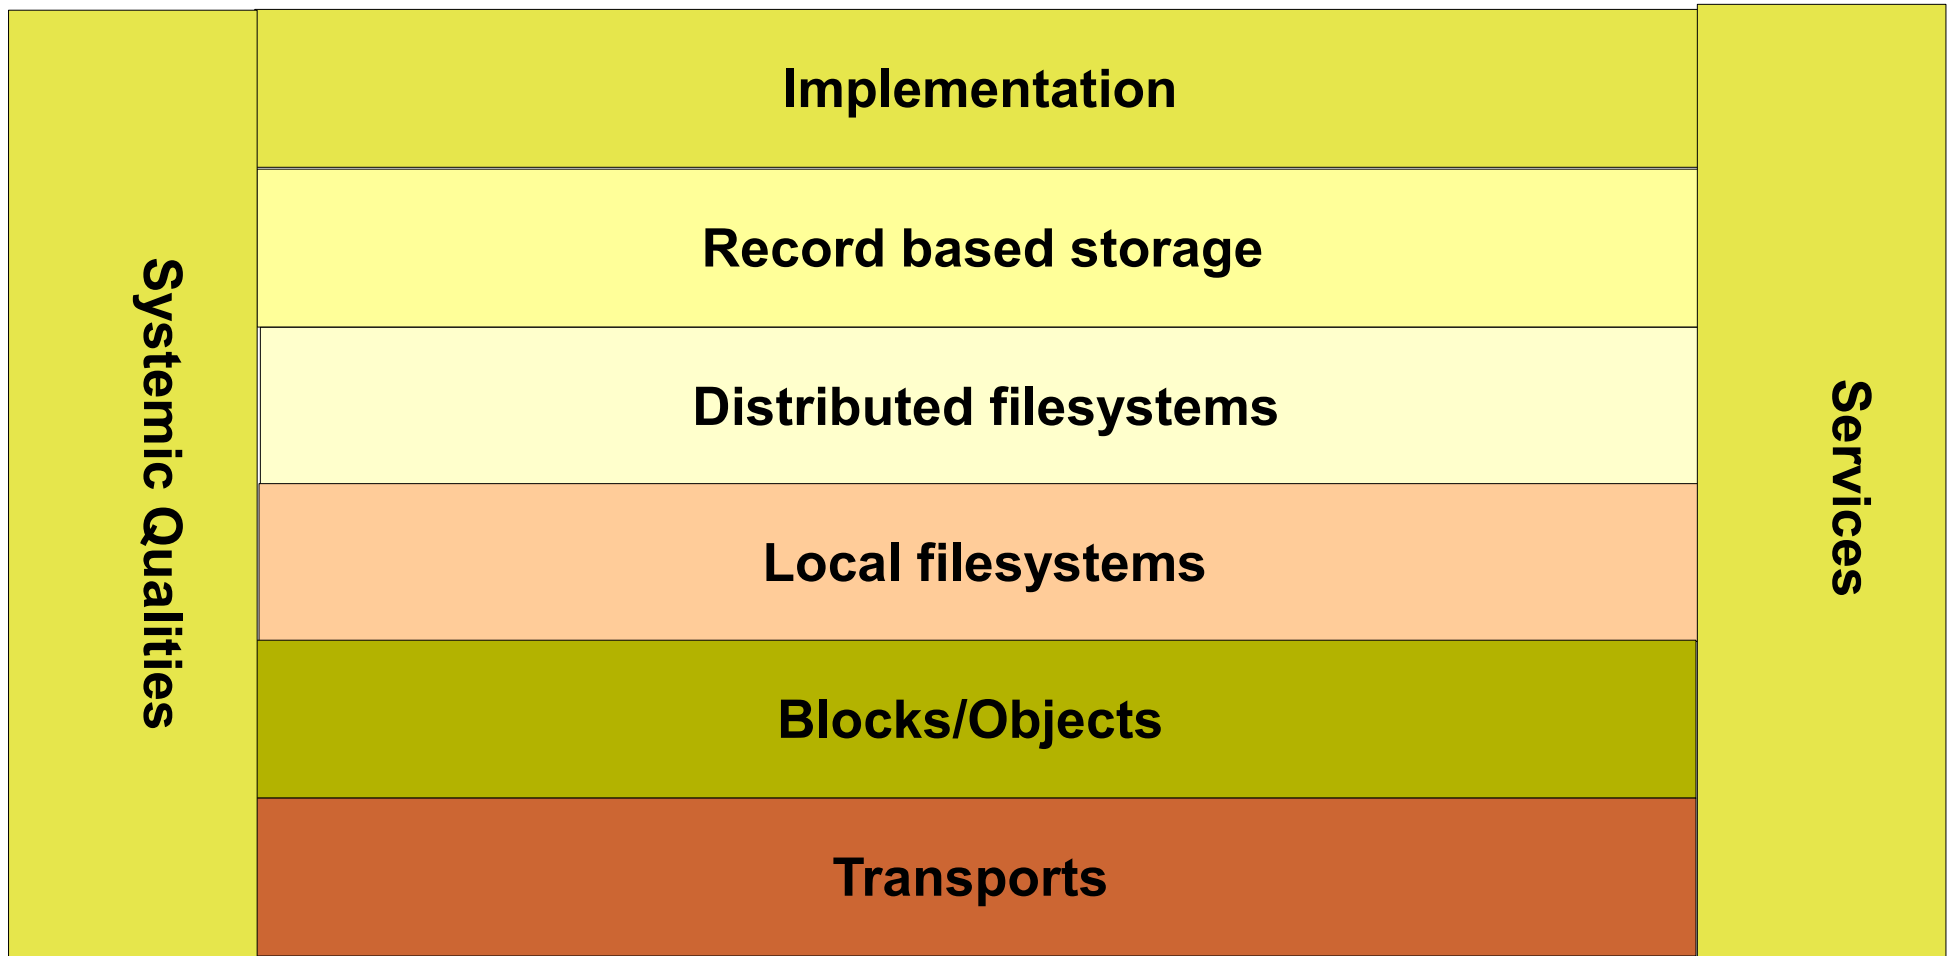


Technology Trends

- General Purpose Hardware
 - Data Servers, Quad Core and Multi-Threaded Servers, Enterprise Class Entry Servers, High Capacity Disks
- Virtualization and Provisioning
 - Server, Network, Storage
- Power Efficient Storage and Servers
- Increased Security and S/W Managed Data Integrity
- Adoption of SAS and SATA
- Adoption of iSCSI over Ethernet



Open Storage Architectural Model



Hardware

Open Storage: Implementation

Availability, Useability, Scaleability, Performance, Security, Flexibility, Cost	MySQL, HADB, Postgres, ...	Block aggregation (RAID), compression, encryption, namespace, discovery, caching, access control, configuration, HSM, HA, replication, snapshot, ...
	NFS, pNFS, CIFS, Shared QFS, Lustre, Honeycomb	
	ZFS, UFS, QFS, SAM-FS	
	SVM Volume management, AVS, MPxIO, iSNS service providers, iSCSI, OSD, FCoE software targets/initiators SD, ST, SES, SCSI, SAS, ... driver framework	
	Fibre Channel, Ethernet, Infiniband, ... driver framework	

Hardware



Systemic Qualities

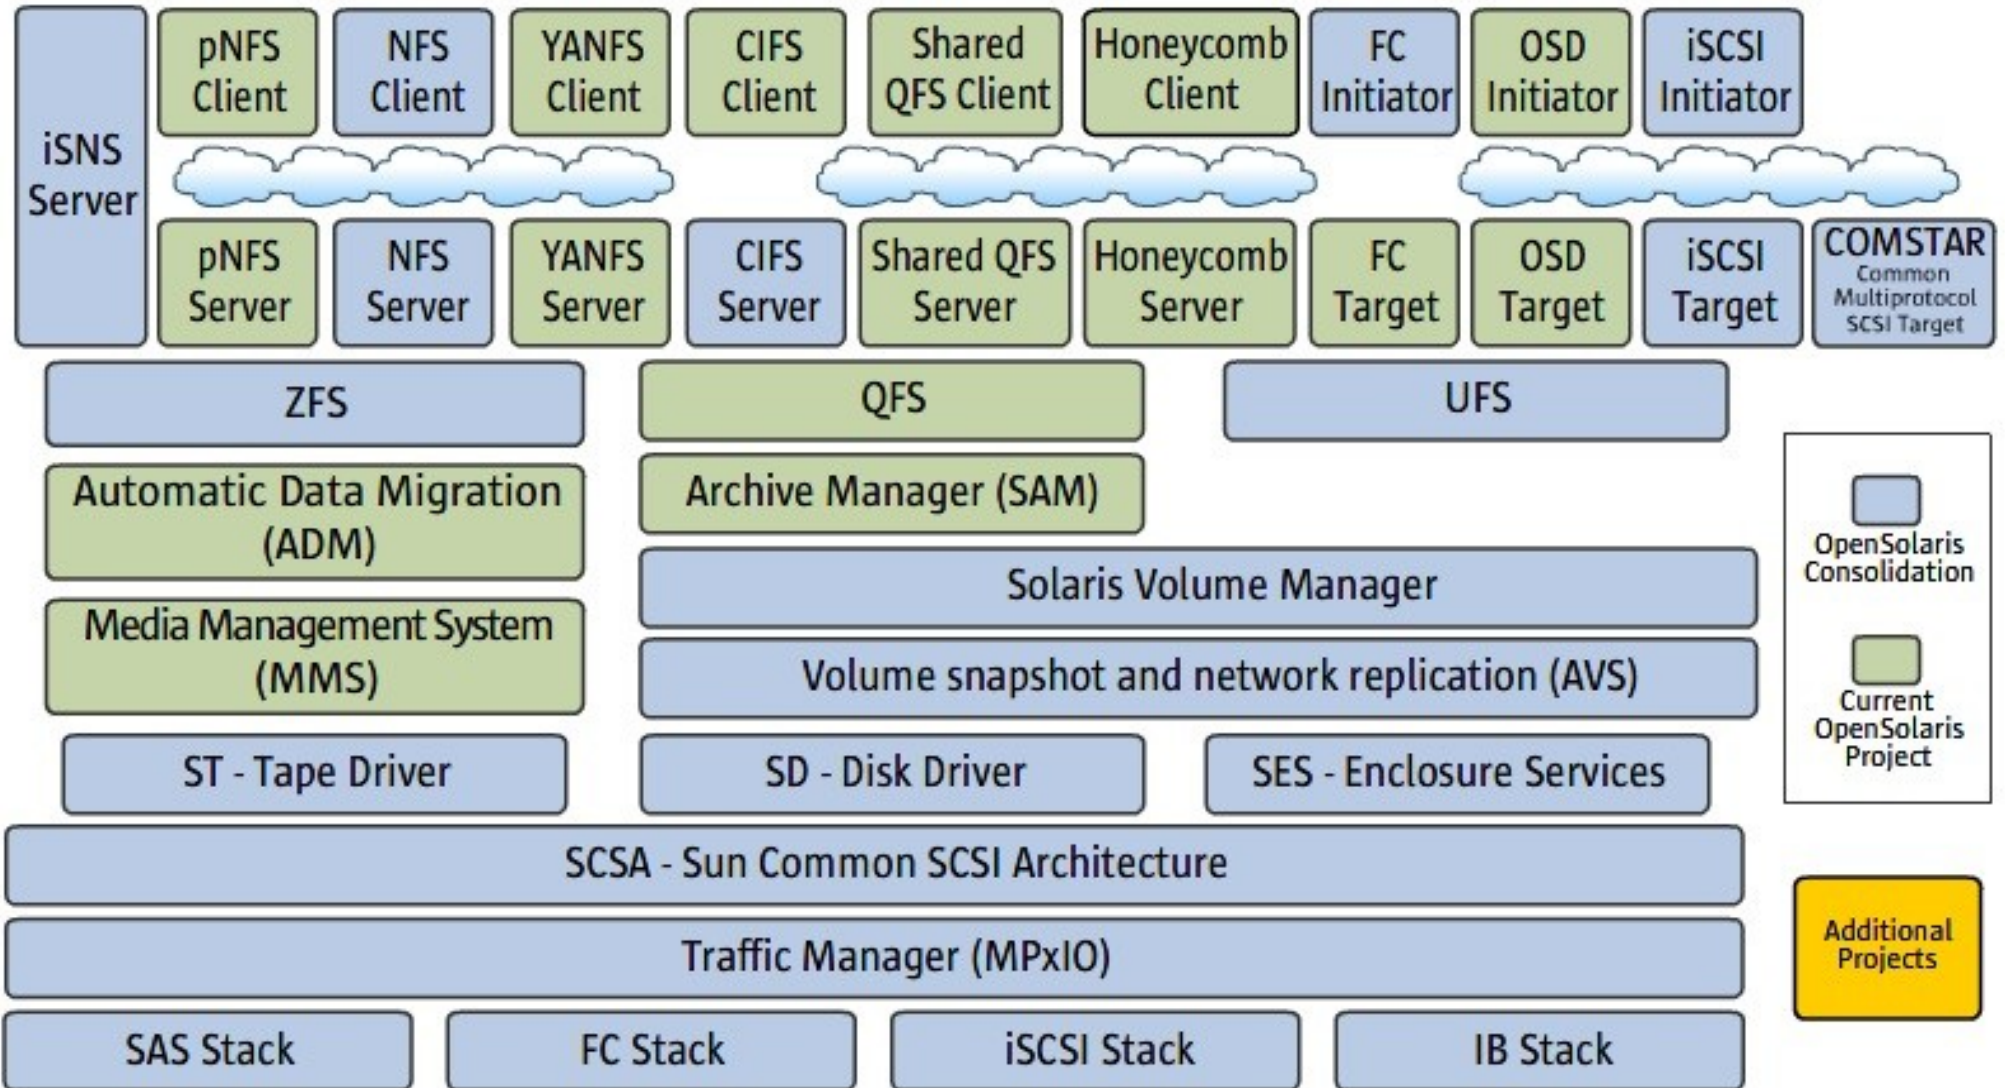
- Reliability, Availability, Serviceability
- Flexibility, Scaleability & Performance
- Security
-
- <your favourite systemic quality>
-
- Cost



Data Services

- Block Aggregation (RAID)
- Replication
- Snapshot
- High Availability
- Policy Based Data Classification
- Encryption
- Data Sharing
- Access Control
- Namespace Management
- HSM
- Device Discovery
- Caching
- Multipathing
- Data Integrity
- Compression
- Archiving
- QOS


Solaris Storage Software





Open Storage Solution Example

Multi-Protocol Storage Access Example

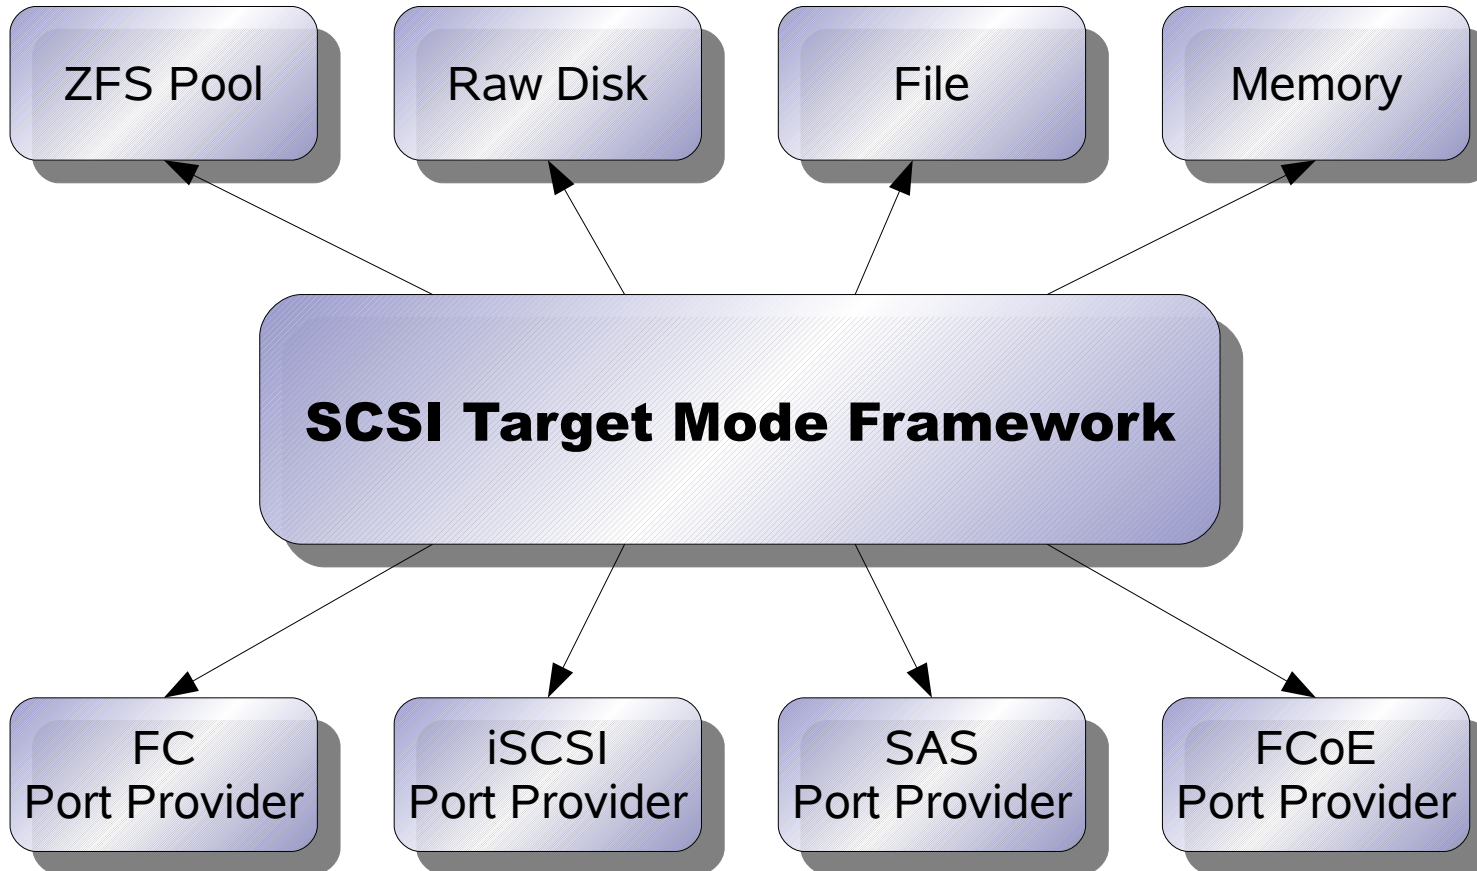


Solaris-based General Purpose Storage

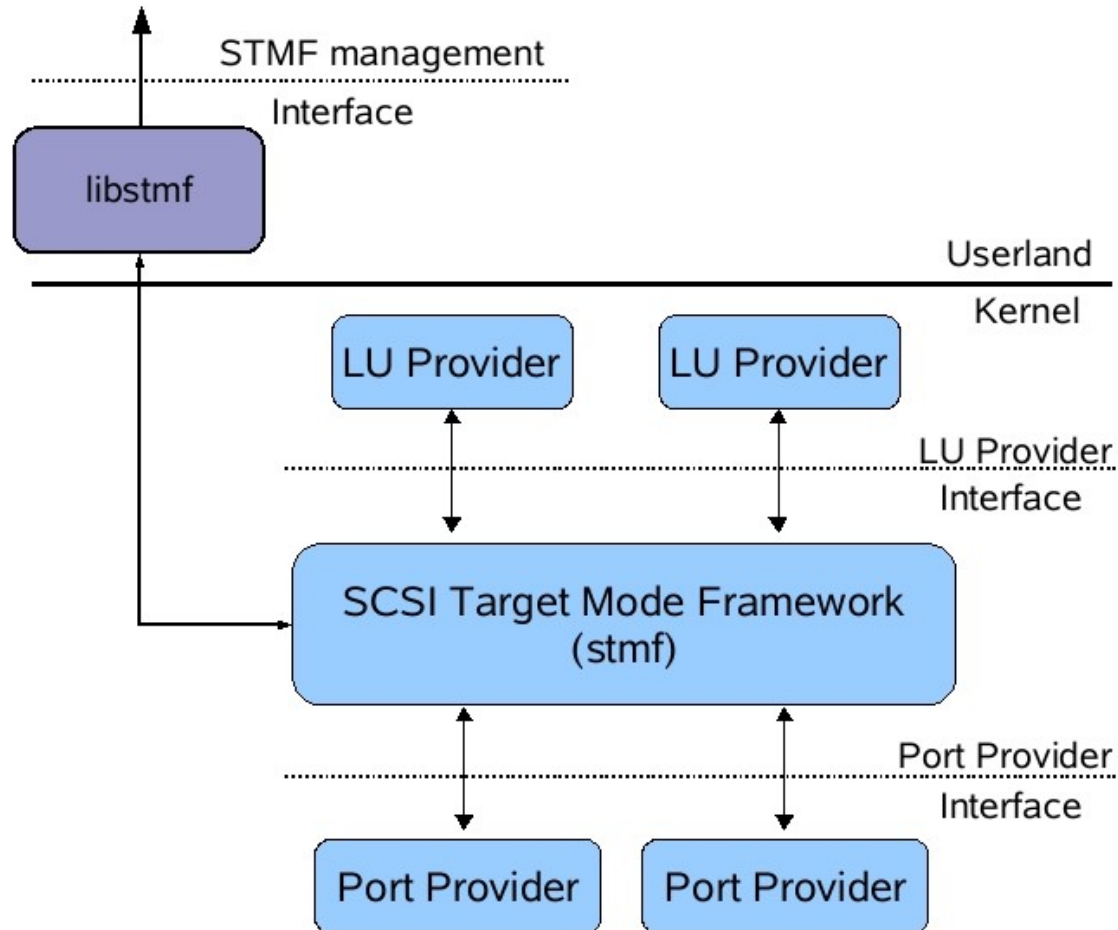


- This is a 4-way x64 server.
- Up to 48 SATA HDDs in 4U.
- It is Sun Fire X4500.
- It can also become:
 - An FC disk array.
 - A NAS box.
 - A SAS array.
 - An FCoE array.
 - How?
 - Solaris with COMSTAR!

Common Multi-protocol SCSI Target



Project COMSTAR (Common Multiprotocol SCSI Target)





ZFS as backing store for COMSTAR

- A file system for everything from desktops to data centers.
- The only self-healing, self-managing general purpose file system.
- Get Into the Pool.
- Delivering Near-Zero Administration.
- Taking the Guesswork Out of Data Integrity.
- Creating Immense Capacity.

COMSTAR: iSCSI Target

- **1. Target**

- Get the service:
- `# svcadm enable stmf ; stmf disable iscsitgt`
- `# itadm create-target ; itadm list-target` # Create targ with iscsi target driver
- `# mkfile 100m /lun ; sbdadm create-lu /lun` # Create our LUN with block driver
- `# sbdadm list-lu (get No.) ;`
- `# stmfadm add view No.` # Enable access to initiator via Comstar
- `# stmfadm list-lu ; itadm list-target (Get No.)`

- **2. Initiator:**

- `# iscsiadm add static-config No,IPaddr:3260`
- `# iscsiadm modify-discovery -s enable`
- `# iscsiadm list static-config`
- `# devfsadm ; format`

On same server - add NFS + CIFS

```
# zpool create -f mypool <LUNS>
# zfs create mypool/myfs
# svcadm enable nfs/server
# zfs set sharenfs=on mypool/myfs
# zfs create -o casesensitivity=mixed mypool/myfs2
# svcadm enable smb/server
# zfs set sharesmb=on mypool/myfs2
# sharemgr show -vp
default nfs=()
zfs zfs/mypool/myfs
nfs=()/mypool/myfs zfs/mypool/myfs2
smb=() mypool_myfs2=/mypool/myfs2
```



Storage Links:

Network Storage:

<http://www.opensolaris.org/os/project/nws/>

COMSTAR:

<http://www.opensolaris.org/os/project/comstar/>

ISER:

<http://www.opensolaris.org/os/project/iser/>

FCoE:

<http://www.opensolaris.org/os/project/fcoe/>

OSD:

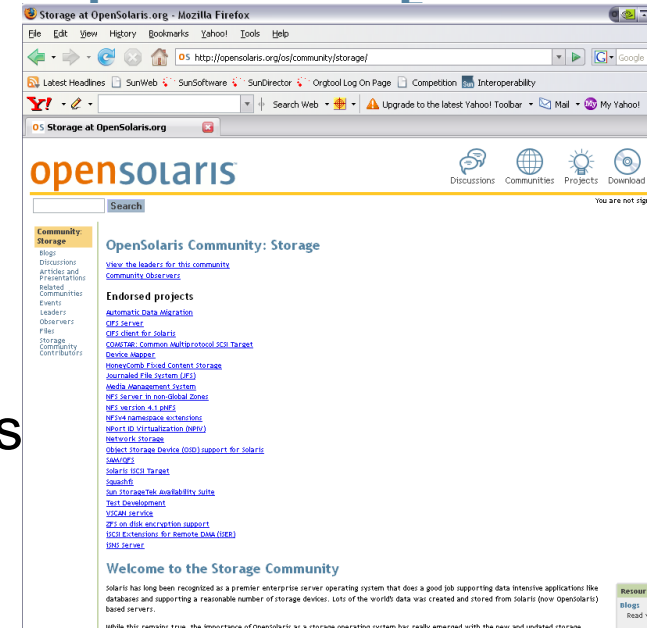
<http://www.opensolaris.org/os/project/osd/>

ZFS:

<http://www.opensolaris.org/os/community/zfs/>

Open Storage Community

opensolaris.org/os/community/storage



- What is the current status?
 - Over 1,500 members (90,000 for opensolaris)
 - 23 endorsed projects
- Who are the members?
 - device developers, appliance developers, storage app. developers, storage adms, system adms
- How do they engage Sun?
 - Mostly forums
 - BigAdmin
 - Sun Developer Network articles
- How do you leverage this community?
 - Code contribution
 - Project feedback
 - Article contribution





Recently Opening Projects

- Celeste
 - <http://research.sun.com/projects/dashboard.php?id=106>
- CAM
- Pegasus
- DAVFS
- MegaSAS



Call to Action

- Evaluate Open Storage Architecture.
- Participate in the Community.
- Open Storage Architecture Resource wikis.sun.com/display/openstorage.



Participate in the Revolution

- **Get Informed!**
 - Watch community town hall, download white papers at:
 - www.sun.com/openstorage
 - Think ZFS:
 - http://www.sun.com/software/solaris/zfs_learning_center.jsp
- **Get Involved!**
 - Join the OpenSolaris Storage community:
 - <http://www.opensolaris.org/os/community/storage/>
 - Join the discussion:
 - <http://www.facebook.com/group.php?gid=12774638>

open



USE



IMPROV
E



EVANGELIZ
E

Thank you!

Dominic Kay
Solaris Marketing
dominic.kay@sun.com
blogs.sun.com/dom

“open” artwork and icons by chandan:
<http://blogs.sun.com/chandan>

開
放
的
열린
مفتوح
libre
मुक्त
ಮುಕ್ತ
livre
libero
ముక్త
开放的
açık
open
nyílt
•••••
πικρ
オープン
livre
ανοικτό
offen
otevřený
öppen
открытый
வெளிப்படை