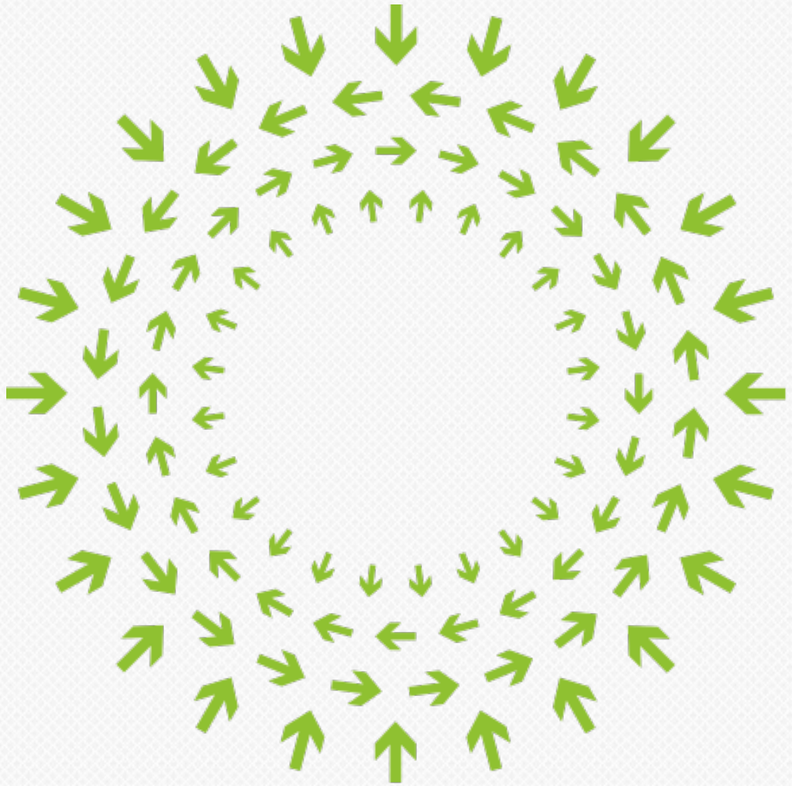

OPEN Compute Summit

Engineering Workshop

October 30-31, 2014

Paris

OPEN Compute Summit

State of the Storage

OCP Storage Committee Report

Asghar Riahi

OCP Storage Chairman

Principal Cloud Architect - Seagate

Agenda

- Introduction
- Personal Story
- OCP Storage Committee Activities
  - Approved Contributions
  - Contributions Currently Under Review
- OCP Storage Call To Action
- OCP Storage Resources
- Links

Introduction Asghar Riahi / Seagate

- Chairman of the OCP Storage Committee since Oct. 2013 - First elected committee chairman in OCP

- asghar.riahi@ocproject.net, asghar.riahi@seagate.com

Background:

- Principal Cloud Architect, Seagate / Cupertino 2 Years
- Master Technologist , HP / Paolo Alto, CA 13 Years
- Unix Systems Manager, MCI Systemhouse Data Center / Napa, CA 2 Years
- System programmer, Siemens AG / Austria & Germany 7 Years
- MS Computer Science / Vienna University of Technology

A Short Personal Story

comp.os.minix › Hello everybody out there using minix – August 25 1991

I'm doing a (free) operating system (just a hobby, won't be big and professional like gnu) for 386(486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhat (same physical layout of the file-system (due to practical reasons) among other things).

I've currently ported bash(1.08) and gcc(1.40), and things seem to work. This implies that I'll get something practical within a few months, and I'd like to know what features most people would want. Any suggestions are welcome, but I won't promise I'll implement them :-)

Linus (torv...@kruuna.helsinki.fi)

PS. Yes - it's free of any minix code, and it has a multi-threaded fs. It is NOT protable (uses 386 task switching etc), and it probably never will support anything other than AT-harddisks, as that's all I have :-(.

https://groups.google.com/forum/#!msg/comp.os.minix/dlNtH7RRrGA/SwRavCzVE7gJ

In 1991 I started working on my master thesis with the topic: ISO-OSI 7 Layer Simulation at Vienna University of Technology and needed to borrow some books. Then I saw a message….

OCP Storage Committee Activities

Approved Contributions

OCP Storage Project Activities

Reviewed and approved four contributed projects since Oct 2013 as follow:

- Storage device with Ethernet Interface — Seagate
- Decathlete Server Board Standard v1 — Intel
- Open Vault Storage Hardware V0.8 — Facebook
- Cold Storage Hardware v0.7 — Facebook

Storage device with Ethernet interface by Seagate

OPEN Compute Summit

Storage device connector

The storage device with Ethernet interface plug connector is the Device Free (Plug) connector defined in SFF-8482 and SFF-8680. See the SFF specifications for detailed dimensional requirements.

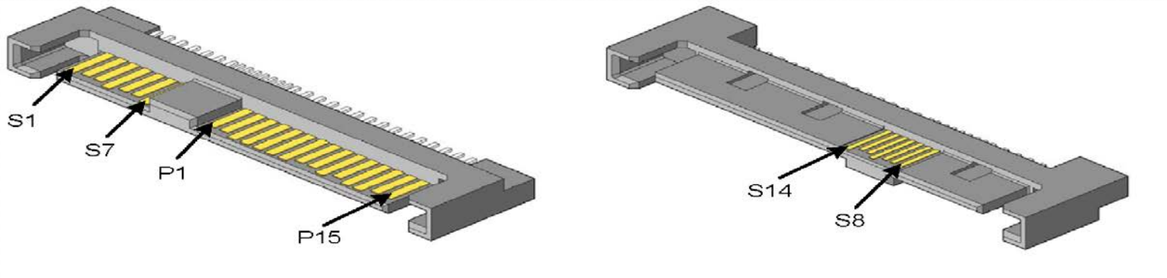

System connector

The system backplane receptacle connector is the Backplane Fixed (Receptacle) connector defined in SFF-8482 and SFF-8680. The backplane receptacle connector defined by SFF-8639 MAY also be used.

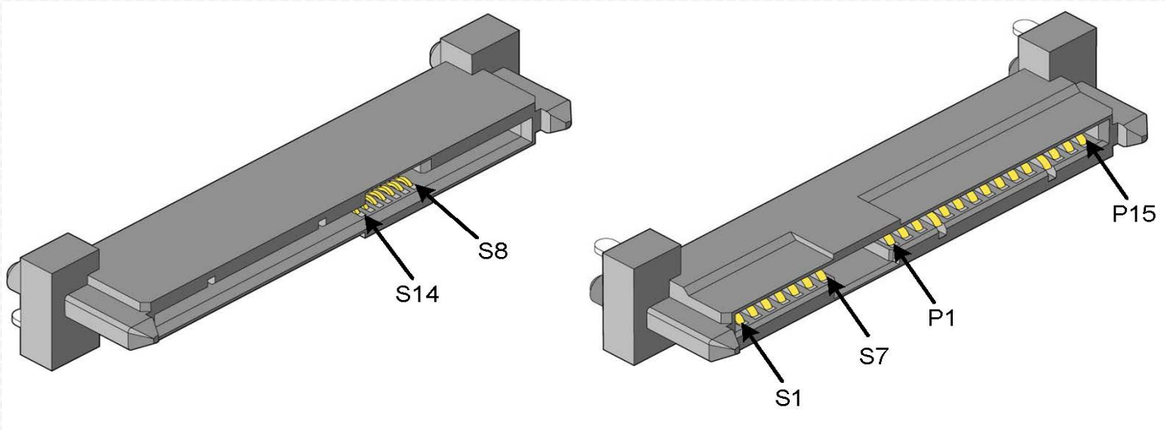

Seagate

Storage device with Ethernet interface by Seagate

OPEN Compute Summit

Single port T-Card for storage device with Ethernet interface

T-Card storage device receptacle connector

The single port T-Card card uses a Backplane Fixed (Receptacle) connector defined in SFF-8482 and SFF-8680 to connect to a storage device with an Ethernet interface. The example single port T-Card in this specification uses straddle mount version of the Fixed (Receptacle) connector.

Seagate

Storage device with Ethernet interface by Seagate

OPEN Compute Summit

2.5” hard disk drive with Ethernet interface plug connector

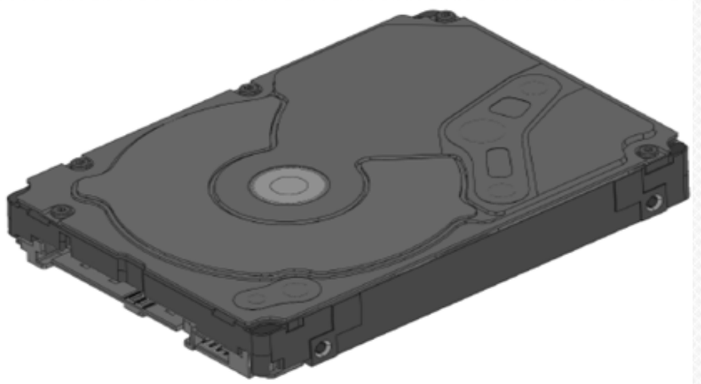

Device form factor and connector location

The storage device with Ethernet interface form factor SHALL comply with SFF-8201 or SFF-8301 (2.5” and 3.5” drive form factors, respectively).

3.5” hard disk drive with Ethernet interface plug connector

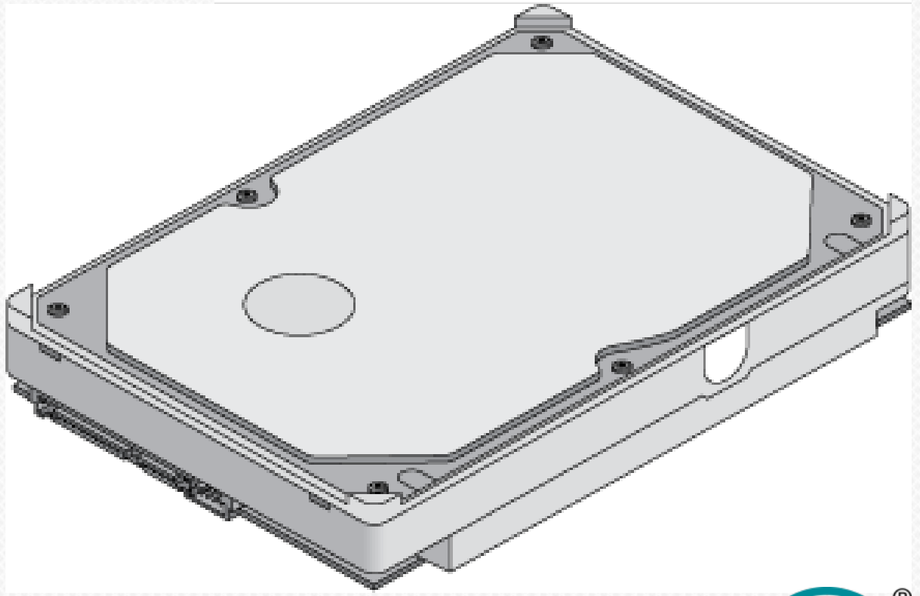

Seagate

Storage device with Ethernet interface by Seagate – Use cases

OPEN Compute Summit

BigFoot Storage Object - The Storage System Revolution!

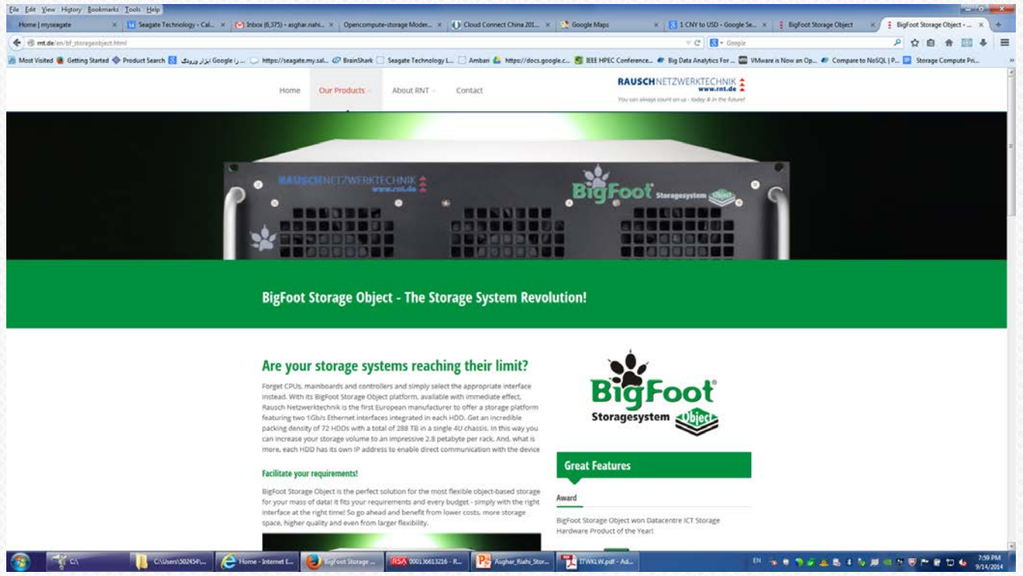

Key features:

- Object based storage
- Up to 288TB capacity in only 4U
- Supports Seagate Kinetic HDDs
- Each HDD with 2x 1Gb/s LAN interface
- No SAS/SATA-controller required
- Only 750mm (29,5") depth for 1,000mm (40") racks

RAUSCH NETZWERKTECHNIK
www.rnt.de
Sympathisch und gut beraten. Bestens betreut.

+ Seagate

Links

What does 288 Terabytes of non-SAS or SATA storage get you?

http://cloud.media.seagate.com/2014/05/20/big-foot-object-storage-storage-hardware-product-of-the-year/

In Chinese Language:

http://www.seagate.com/cn/zh/case-studies/rausch-bigfoot-scale-out-object-storage-seagate-kinetic-platform-cs/

Rausch Systems:

http://rnt.de/en/bf_storageobject.html

Seagate Kinetic Solution: A look at the Rausch BigFoot Object Storage solution

http://www.techradar.com/news/computing-components/storage/seagate-kinetic-solution-a-look-at-the-rausch-bigfoot-object-storage-solution-1232657

Open Vault Storage System Using Ethernet Storage Device NEB instead of SEB

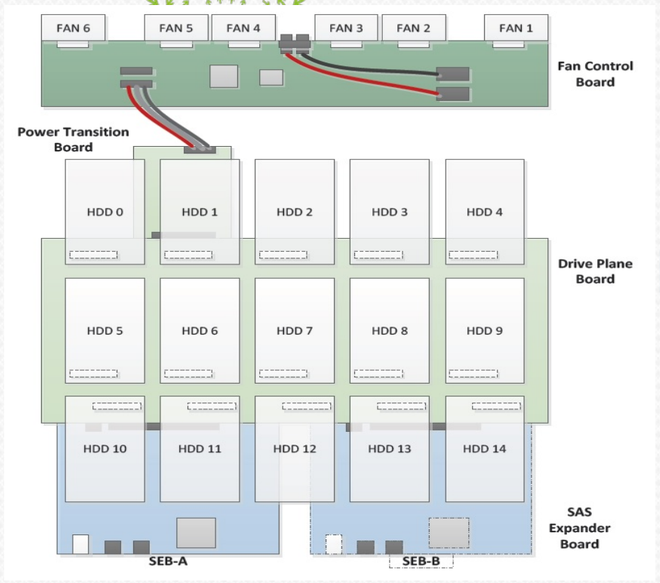

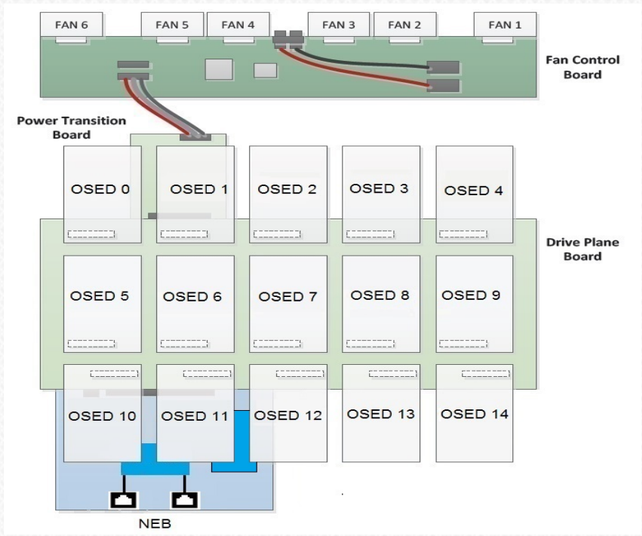

8 Knox storage slots per one Winterfell server = 16 *30 = 480 HDD per Rack

No more Storage nodes needed 540 Storage Devices 18 *30 = 540).

Decathlete Server Board Standard by Intel

Decathlete Server Board Features

The Decathlete Server Board is intended to meet the most common usages for 1U and 2U dual-socket servers in the scalable data center.
To insure the delivery of products that can be deployed over a period of time, and assure consistency in the services offered to the client of the cloud data server, certain features must be present in each model or generation of servers.

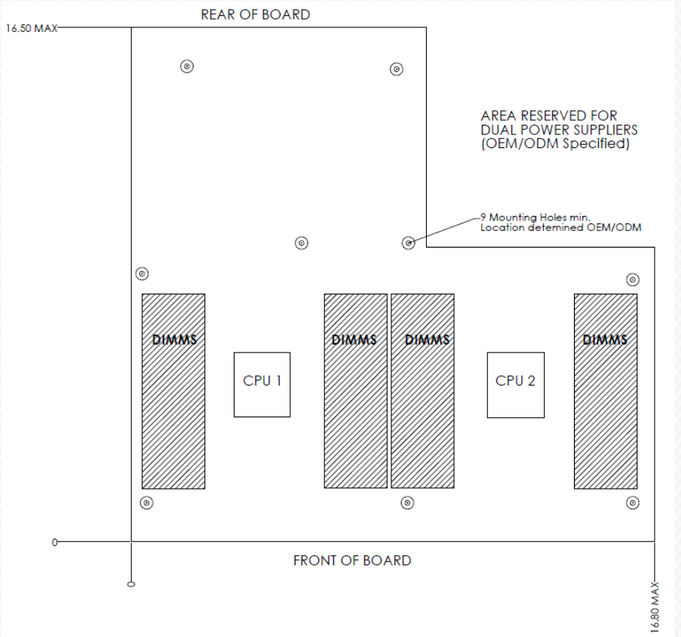

Open Vault Storage Hardware V0.8 by Facebook

OPEN Compute Summit

The Open Vault storage unit is a 2U-30HDD storage enclosure, consisting of two identical 1U high HDD trays with 15 x 3.5" HDDs and slots for two SAS expander boards on each, one fan control board, and six redundant fan modules mounted externally in the rear of the chassis. An Open Vault storage server fits into the Open Compute Project Open Rack.

facebook

Cold Storage Hardware v0.7 by Facebook

A Cold Storage system design comprises, but is not limited to, the following aspects:

- Ability to adopt current and future HDD technologies with the lowest cost
- Capability to power off HDDs that are not in use
- Modification of storage unit (based on Open Vault)
- Configuration of an OCP compute node
- Mini-SAS fan-out cable between the Open Vault and the OCP compute node
- Custom Open Rack for the configuration of the Cold Storage system
- Redefined topology for networking switch deployment
- New power consumption provisioning, and new data center floor plan, and so forth.

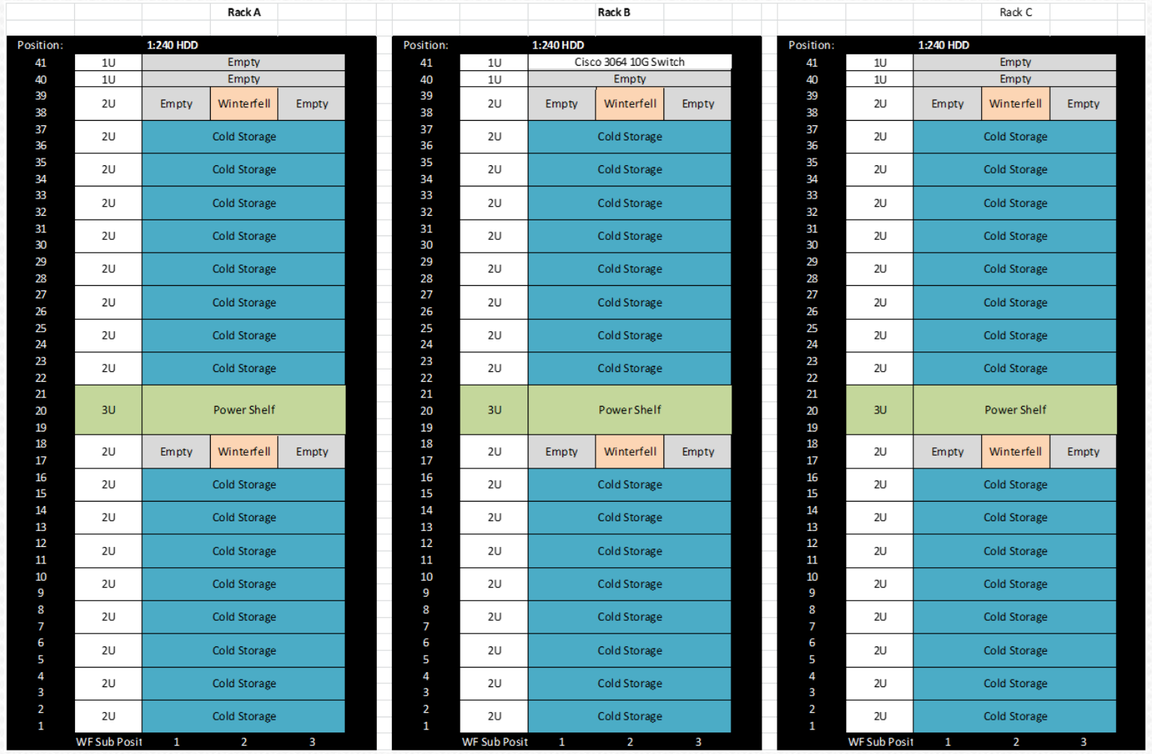

OCP Storage Committee Activities

OPEN Compute Summit

Contributions Currently Under Review

OCP Storage Committee – Projects Under Review

Open Vault Storage Hardware V0.85 Avago

Nytro XP6209 Application Acceleration Card Seagate

OCP C&I Storage Certification Specification and Test Plan OCP C&I

Open Vault Storage Hardware V0.85 By Avago Technologies (Ex. LSI)

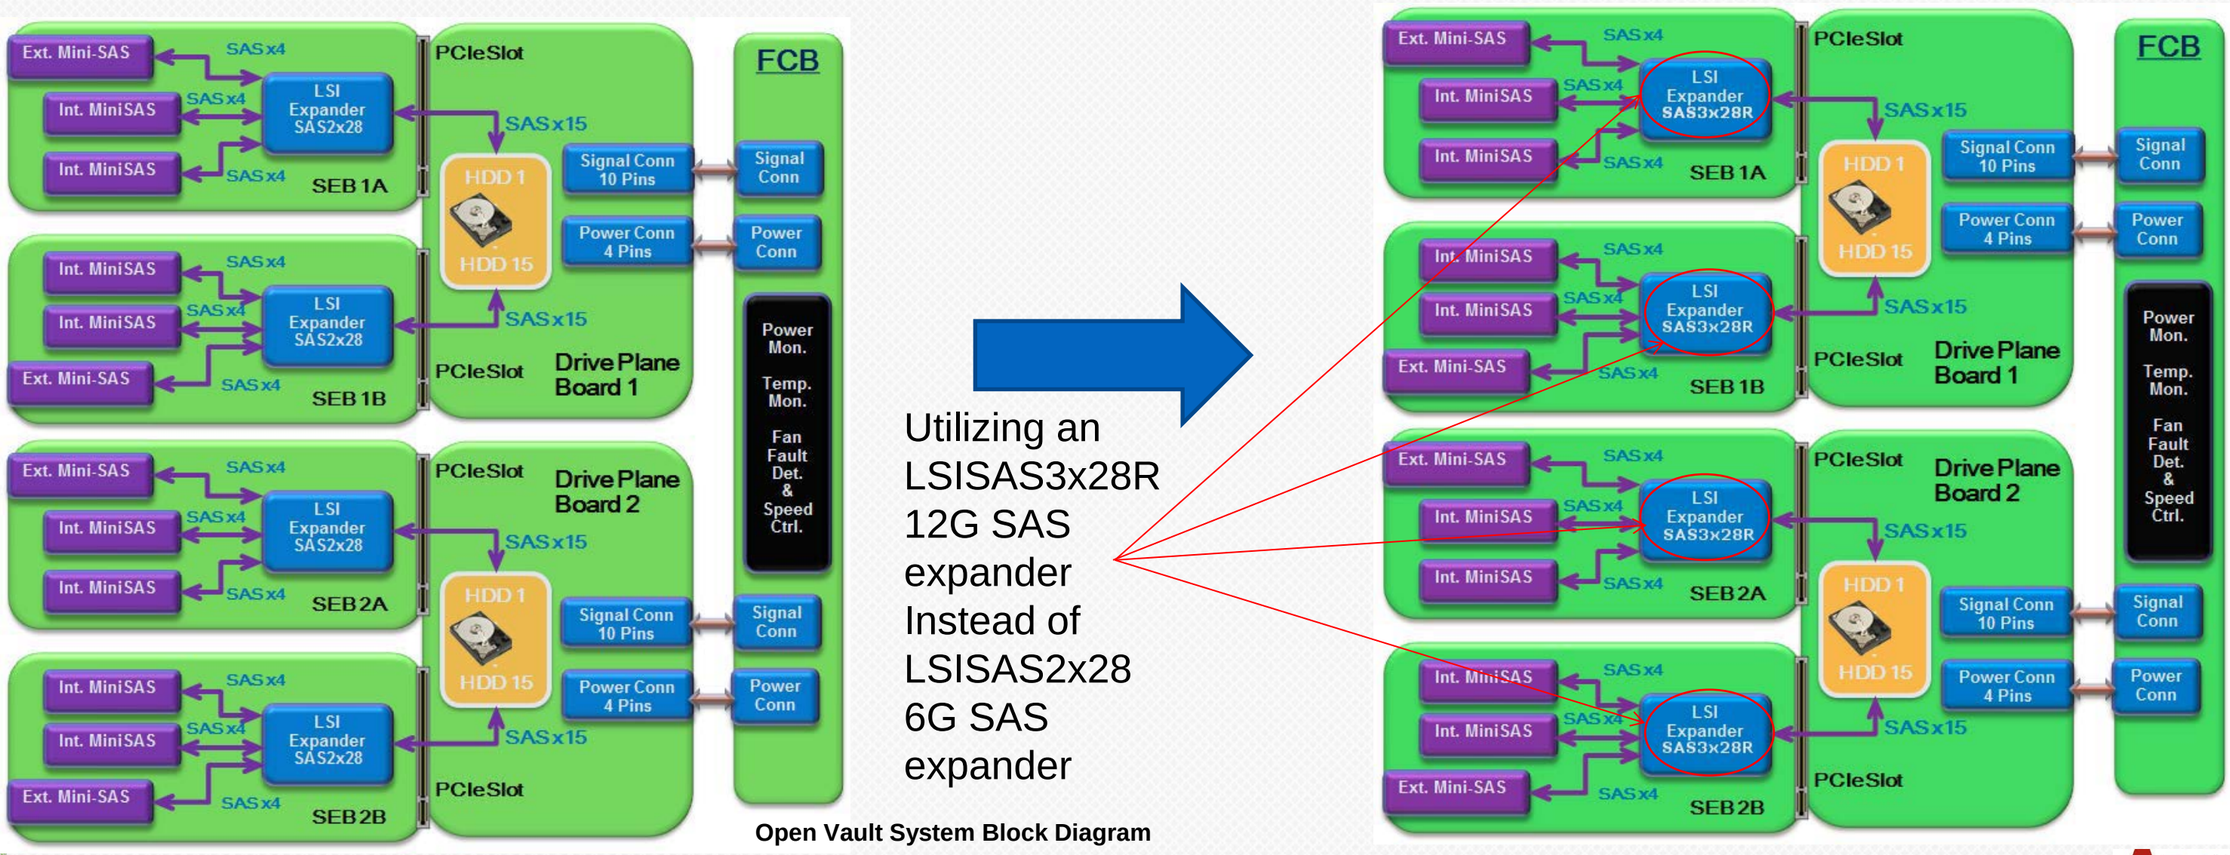

Open Vault System Block Diagram

Nyto XP6209 Application Acceleration Card by Seagate

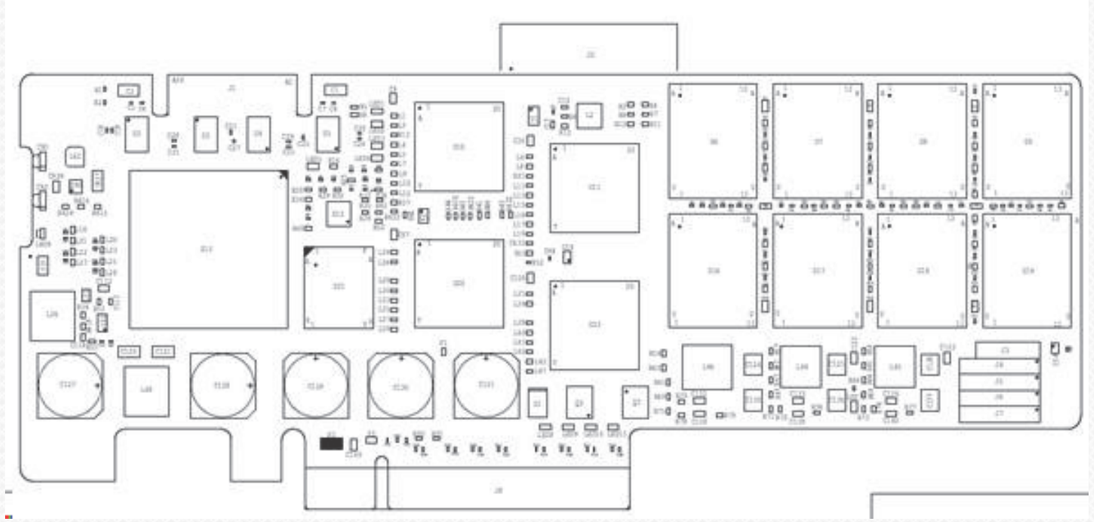

Features

The **LSI Nytro** XP6209 Application Acceleration Card acts as a PCIe-based block storage device and presents itself to the operating system (OS) through a **Fusion-MPT™** interface.

OCP C&I Storage Certification Test Document

OCP C&I Storage
Certification Specification and Test Plan
Version 0.13

The aim of this document is to provide all the information needed to perform OCP Certified testing on a proposed Open Compute compliant platform. It will provide information on getting and installing the testing tools as well as outline several test cases and provide additional information.

UTSA OCP Certification & Solution Laboratory

OCP Storage – Contribute

“Our first step for the Open Compute Project is releasing the specifications and mechanical drawings. The second step is working with the community to improve them” From OCP Sepcs & Designs

“We believe that openly sharing ideas, specifications and other intellectual property is the key to maximizing innovation and reducing operational complexity in the scalable computing space. The Open Compute Project Foundation provides a structure in which individuals and organizations can share their intellectual property with Open Compute Projects. “ – From OCP MISSION STATEMENT

OCP Storage – CALL TO ACTION:

OPEN Compute Summit

OCP Storage Resources

OCP Main Website: http://www.opencompute.org

OCP Storage web site: http://www.opencompute.org/projects/storage/

OCP Storage Wiki: http://www.opencompute.org/wiki/Storage

OCP Storage Project Specs:
http://www.opencompute.org/wiki/Storage/Dev

Email: asghar.riahi@ocproject.net, asghar.riahi@seagate.com

How to Join the OCP Storage

- Check Amber's Get Involved site:
  - http://www.opencompute.org/community/get-involved/
  - Mailing List
  - http://lists.opencompute.org/mailman/listinfo/opencompute-storage
  - Monthly Calls
  - http://www.opencompute.org/wiki/Storage
  - Contact via Email:
  - opencompute-storage@lists.opencompute.org
  - asghar.riahi@ocproject.net, asghar.riahi@seagate.com
  - Meetup.com
  - http://www.meetup.com/Open-Compute-Project/

This is an interstitial slide

THANK YOU

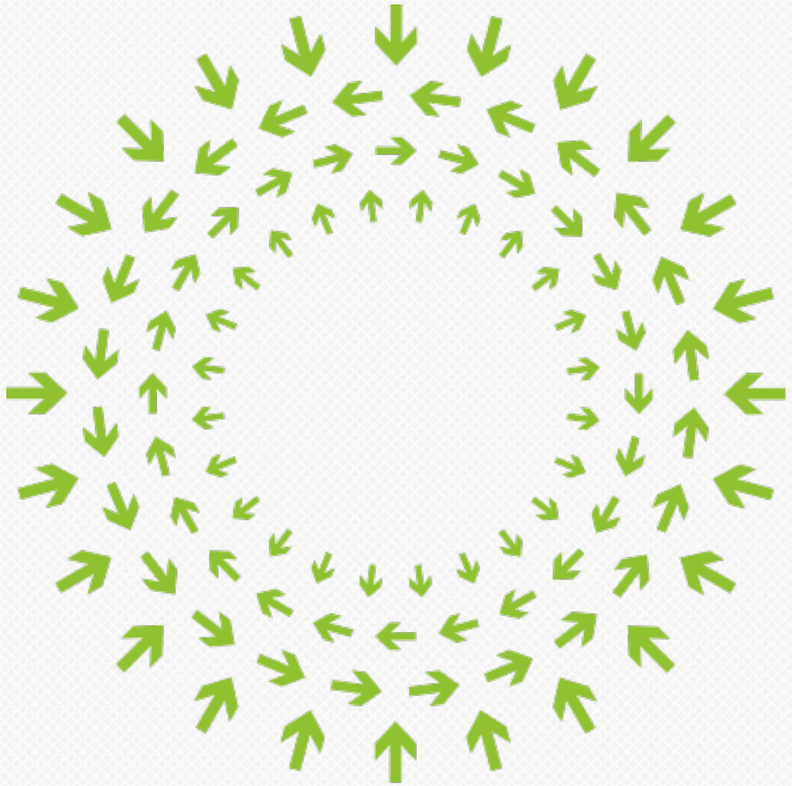

OPEN Compute Summit

Engineering Workshop

October 30-31, 2014

Paris