# Open Source Embedded Cloud Storage

Kai Blin
SAMBA Team

# Outline



- Intro
- Design
- Implementation
- Conclusions

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Intro



- About myself
- Defining "private storage cloud"
- Embedded cloud storage trade-offs

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# About myself



- M.Sc. in computational biology
- Microbiology grad student
- Open Source developer
- Samba Team member

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Defining "Private Cloud Storage"



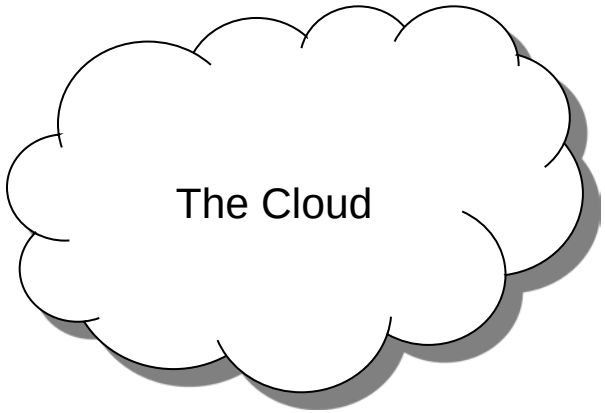


Third Party

**VS.**

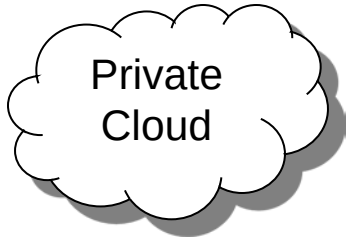


Internal

- Someone else's hardware
- Pay for actually required use
- Need to trust third party

- Your own hardware
- Pay for estimated peak usage
- Data privacy easier

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

SAMBA

# Embedded Cloud Storage Trade-Offs



- Low power consumption
- Living room compatible
- Extensible

- Low performance
- Low disk space

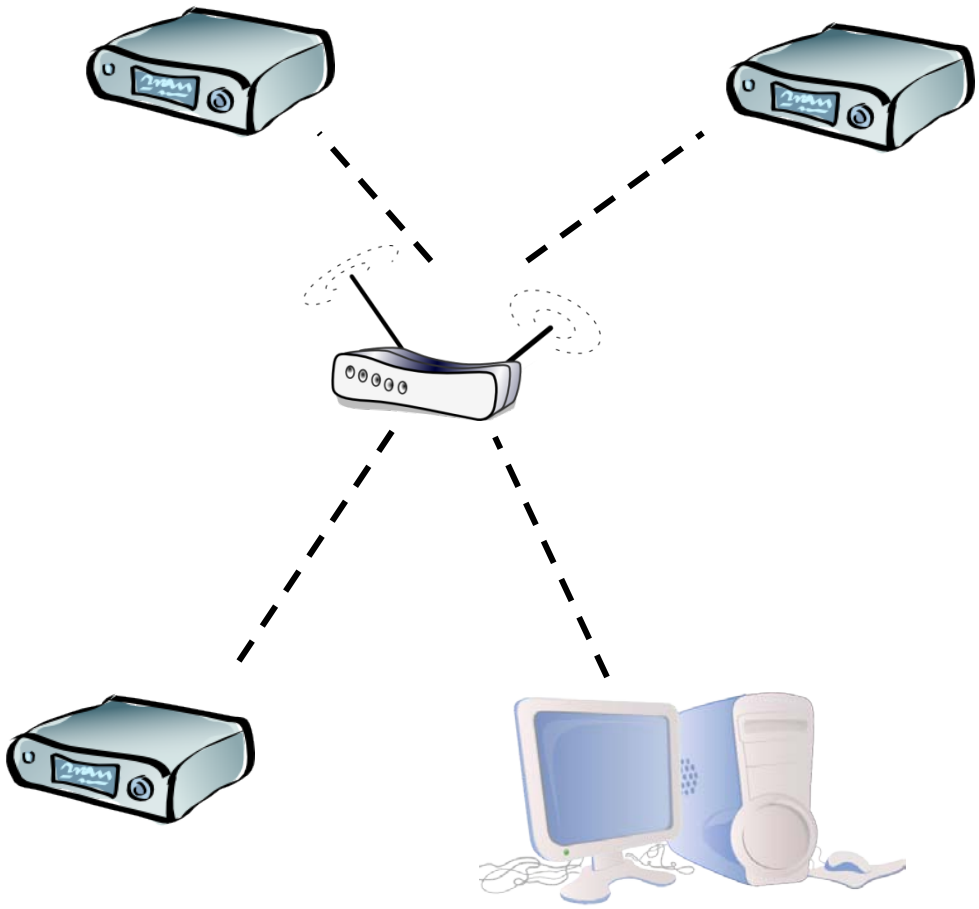

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

SAMBA

# Design



- Goals
- Non-goals
- Hardware Requirements
- Hardware Options
- Software Requirements
- Open Source Cluster File Systems
- Distribution Choices

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Goals



- Cheap storage
- Gadgetable
- Easily expandable
- Low power consumption
- Low noise
- Keeping data for home entertainment or SOHO

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Non-Goals



- Blazingly fast
- Petabyte capacity

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Hardware Requirements



- Fast storage access
- Fast networking
- Lots of RAM
- Lots of CPU power

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Hardware Options



- BeagleBoard
- BeagleBoard xM
- Hawkboard
- Sheevaplug
- Guruplug

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# BeagleBoard



- OMAP3530 @ 720MHz
- 256MB RAM
- USB host port
- USB OTG port
- **NO** ethernet

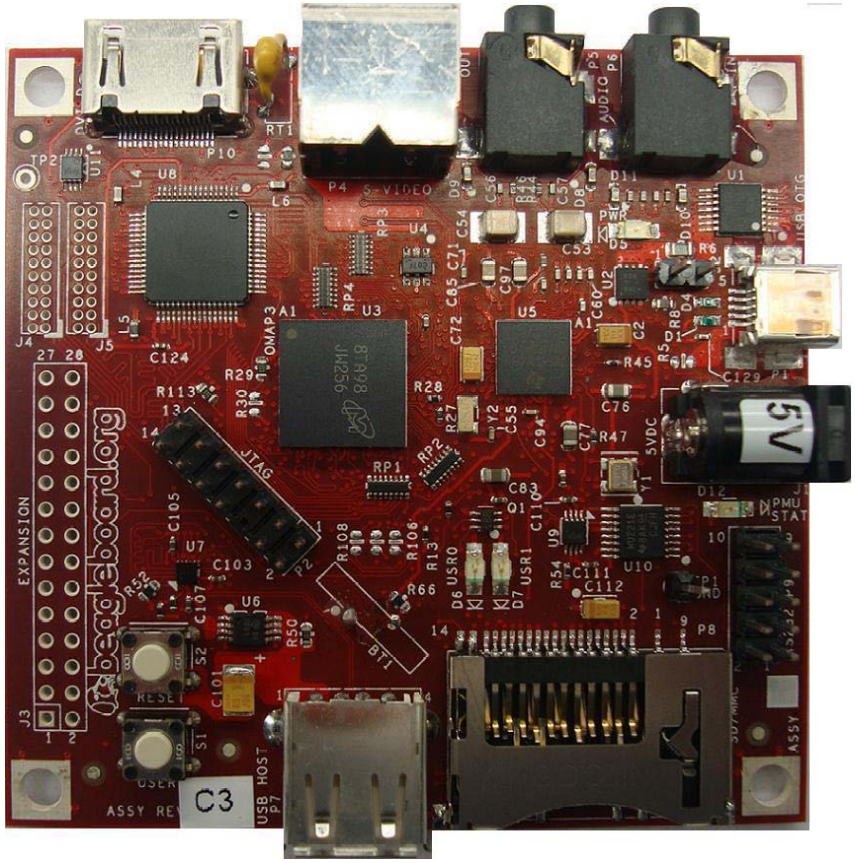


Source: BeagleBoard System Reference Manual Rev C4, Page 1

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# BeagleBoard xM



- DM3730 @ 1GHz
- 512MB RAM
- 4 USB host ports
- USB OTG port
- 10/100 Mbit ethernet

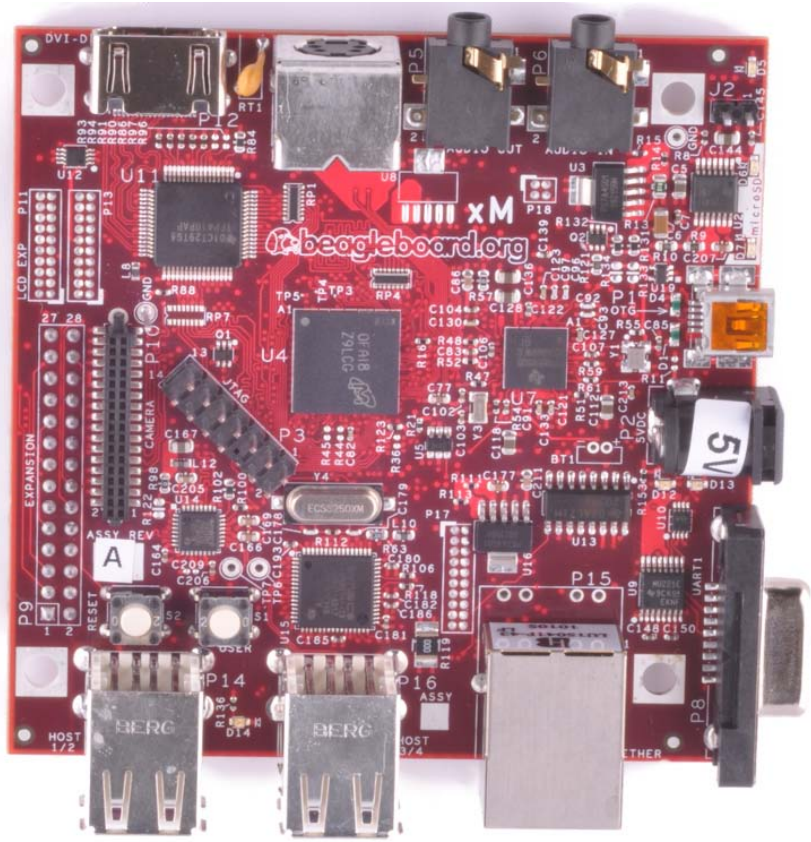


Source: http://www.flickr.com/photos/jadon/4628635196/sizes/o/in/photostream/

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Hawkboard



- OMAP-L138 @ 450 MHz
- 128MB RAM
- USB host port
- USB OTG port
- 10/100 ethernet
- SATA

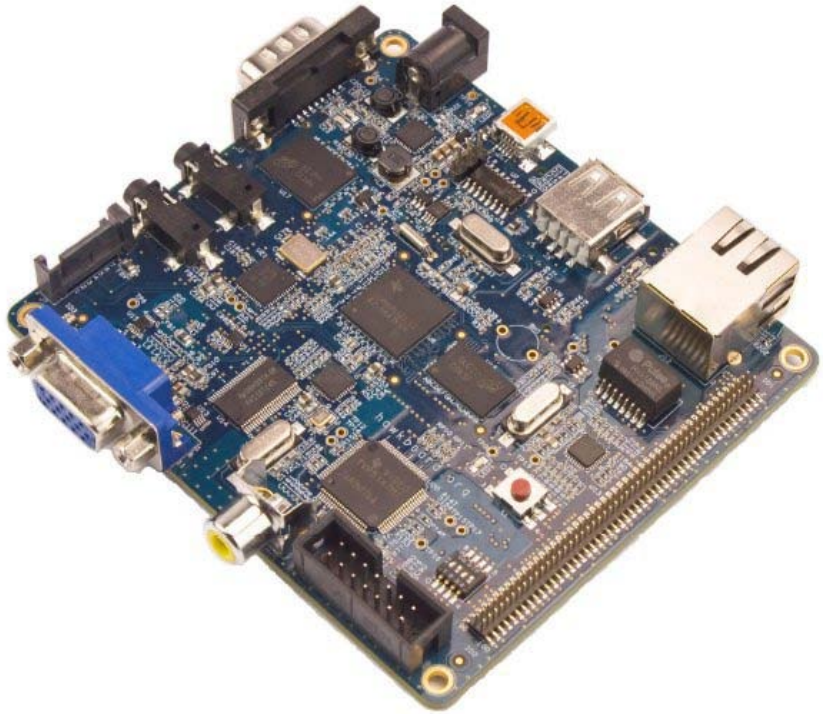

Source: Hawkboard User's Manual, Page 1

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Sheevaplug



- ARM Kirkwood @ 1.2GHz
- 512MB RAM
- USB host port
- 10/100/1000 ethernet
- (some models with eSATA)


Source: http://newit.co.uk/shop/proddetail.php?prod=SheevaPlug

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Guruplug (Server Plus)



- ARM Krikwood @ 1 GHz
- 512MB RAM
- 2 USB host ports
- 2 10/100/1000 ethernet
- eSATA

Source: http://www.globalscaletechnologies.com/t-guruplugdetails.aspx

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Software Requirements



- Open Source platform
- Easy to customize
- POSIXy file system semantics
- Multi-Platform access to data

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Open Source Cluster Filesystems



- Ceph
- Global File System (GFS)
- Gluster
- Lustre
- Moose File System (MooseFS)
- Oracle Cluster File System (OCFS)
- Parallel Virtual File System (PVFS)




2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

SAMBA

# Distribution Choices



- Ångström
  - Focus on GUI-based embedded systems
  - Powerful toolchain
- Ubuntu
  - Widely-used system
  - ARM support focused on netbooks
- Debian
  - Supports older ARM architectures

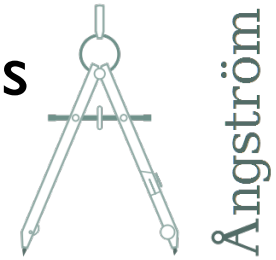


ubuntu®

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

SAMBA

# Implementation



- Hardware used
- Software used
- Cluster layout
- Cluster configuration
- CTDB configuration
- Samba Configuration

Implementation details might still change, get the updated slides after the conference.

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Hardware Used



- Hawkboard
  - SATA-capable
  - 100Mbit ethernet
  - 128MB RAM
- Beagleboard xM
  - USB only
  - 100Mbit ethernet
  - 512MB RAM

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Software Used



- Ångström Linux and Ubuntu Linux
- Ceph, MooseFS and PVFS tested
- CTDB
- Samba 3.6 pCIFS

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Cluster Layout



client

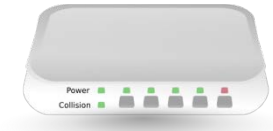


node

node

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

SAMBA

# Cluster Configuration



- Both nodes as object storage devices
- BeagleBoard xM runs other required services
- Both nodes run CTDB
- Both nodes run cluster-aware Samba 3.6
- Laptop as CIFS client

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# CTDB Configuration



- Disable recovery locks (potentially dangerous)
- Set up public interfaces with takeover
- Allow CTDB to manage Samba
- Set up DNS round robin

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Samba Configuration



- Compile Samba with cluster support
- Set `clustering = yes`
- Set `idmap backend = tdb2`
- Turn off mmap support
- Normal Samba configuration otherwise

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Conclusions



- Performance data
- Viability
- Questions & Answers

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Benchmarks



- To be done

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Power Consumption



- To be done

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Viability



- To be done

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Thank you for your attention



Questions?

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.

# Links



- Samba: http://samba.org
- CTDB: http://ctdb.samba.org
- Ceph: http://ceph.newdream.net
- Ångström: http://angstrom-distribution.org
- BeagleBoard: http://beagleboard.org
- Hawkboard: http://hawkboard.org
- Further links: http://kblin.org/talks/sdc2010.php#links

2010 Storage Developer Conference. Kai Blin. All Rights Reserved.