Ark Storage Solutions

James R. Bergsten, Founder
jim@thebergstens.com
April 27, 2013

Executive Summary

Ark Systems owns internally developed Intellectual Property (IP) consisting of:

- ✓ A highly efficient, easy to support, portable, multiprocessing Real Time Operating System (RTOS)
- ✓ A full-featured, data storage appliance application

Ark seeks to be acquired by an established enterprise wishing to leverage Ark's IP to bring new product offerings to the marketplace

ARK SYSTEMS™

Ark History

- Ark *Research* was formed in 1995 to provide IT disaster recovery and business continuance solutions for the midrange open systems marketplace
- Ark was acquired by LSI Logic in 2000 ($23M cash) to give LSI's customers remote data mirroring capabilities
- Ark and LSI installed over sixty systems worldwide that ran for years without a single high severity problem
- Due to business strategy changes, LSI quiesced the project in 2002
- Ark *Systems* was formed in 2005 to spin-off, significantly enhance the IP, and reintroduce storage products to the marketplace
- Ark is funded by its founder – it never reached critical mass and was quiesced in mid-2008, though product development continues

ARK SYSTEMS™

Ark Value

- Cross industry – used "everywhere" – industry, medical, power/energy, telecommunications, IT, manufacturing…
- Scalable – horizontal (multiple nodes) or vertical
- Portable – remove dependency on single platform
- Build customer base – compete against the small, migrate toward the big – lower end cheaper COGS
- Simple – easy to control software process, develop, predictable, lower NRE and OPEX, faster TTM, simpler upgrade
- Robust – high quality, reliability, and provable correctness lead to easier certification
- Build IP resource, introduce data-storage offerings, add expertise
- No GPL / overseas issues
- Capture customers / applications

ARK SYSTEMS™

Architecture

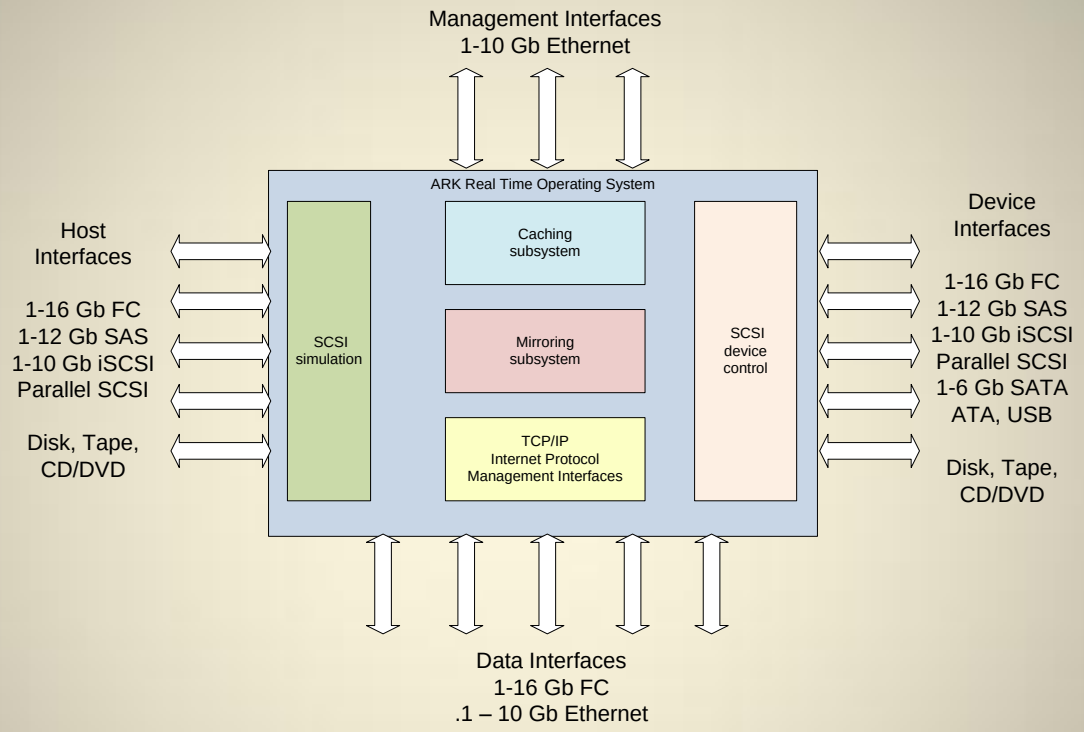

ARK SYSTEMS™

Data Storage Applications

- Heterogeneous, transparent host and device access
- Local and remote block-granular device virtualization
- Large, preemptive cache
- *N-way* volume mirroring, cascade mirroring
- Protocol translation
- Non-disruptive device testing and data migration
- No architectural restrictions – vertically and horizontally scalable
- Twelve issued U.S. Patents
- Runs under Ark RTOS but could be ported to other OS'es
- Small footprint (RTOS + Application approx. 2MB x86)
- Quick, simple build ("compile → link → load → boot" < 30 seconds)

ARK SYSTEMS™

Storage IP Market Value

- Runs on anything – no hardware NRE costs
- Better margins than existing alternatives
- Immediate access to IP, bring storage appliance business in-house
- Connects anything to anything
- The implementation is not vendor-specific
- Mature business continuation and disaster recovery technology
- Data-storage-specific middleware offering
- Ark is a data storage "Swiss Army Knife"

ARK SYSTEMS™

Existing Solutions

Using Present Technology…

- ✓ Bridging (shelf modules)
- ✓ In-stream appliance to add advanced functionality to existing storage
- ✓ Disaster recovery / business continuation
- ✓ Virtualization, consolidation, zoning, partitioning
- ✓ Non-disruptive data migration
- ✓ Solid-state / FLASH disk
- ✓ Vertical and horizontal scalability

- Product Complete
- Documentation Complete

ARK SYSTEMS™

Industries

- ✓ Education -- Drexel University
- ✓ Telecommunications / Cellular -- MobilCom
- ✓ Entertainment / Media / Video -- MTV
- ✓ Back Office / Payroll Ceridian
- ✓ Real Estate -- ECE
- ✓ Manufacturing / Industrial -- Wacker Siltronic, PPG
- ✓ Legal -- First American Title Insurance Company
- ✓ Retail -- Next Retail Ltd.
- ✓ Health Care -- Robert Bosch Krankenhaus
- ✓ Banking / Finance -- Savings Bank Organization, SWIFT
- ✓ Medical -- St. Olav Hospital
- ✓ Transportation -- Siemens Transportation Systems
- ✓ Utilities -- Southern Company / Alabama Power

ARK SYSTEMS™

Customers

- ✓ Drexel University – Data Backup Optimization
- ✓ MobilCom – Data Consolidation and Data Mining
- ✓ MTV – Business Continuation and Disaster Recovery
- ✓ Ceridian – Long Distance Business Continuation in a Heterogeneous Environment
- ✓ ECE – Campus Distance Business Continuation
- ✓ Wacker Siltronic – Business Continuation
- ✓ First American Title Insurance Company – Microsoft Windows Cluster Disaster Recovery
- ✓ Next Retail Ltd. – Business Continuation
- ✓ Robert Bosch Krankenhaus – Health Care Data Availability
- ✓ Savings Bank Organization – Remote Data Access
- ✓ St. Olav Hospital – Disaster Recovery
- ✓ Siemens Transportation Systems – DR
- ✓ Southern Company / Alabama Power – DR

ARK SYSTEMS™

Potential Applications

- **Blade / Storage Appliance products**

  Switches, bridges, protocol converters, data migration, backup offload, tape simulation…

- **Real-time products**

  Data collection, voting, kiosks, set-top boxes, home entertainment, Internet appliances…

- **Professional Services**

  Data Center testing, migration, performance tuning…

ARK SYSTEMS

Small Office BC/DR

"Set-top" remote mirroring appliance for small enterprises or remote offices

- Similar to Cable box
- Customer installs Ark between server and local storage
- No changes to server, storage, OS, software…
- Remote data is consolidated within CoLo partner sites
- Provides BC/DR for smaller shops
- Customer pays for equipment, setup, subscription, PS

ARK SYSTEMS™

Blade Server I/O Subsystem

- Ark runs on more or more “I/O” blades
- Provides interface to the “outside world”
- Virtual Shared Volumes eliminate need for physical system volume(s) on each blade
- Central management facilitates VM’s
- Licensed to Blade Vendors

ARK SYSTEMS™

Data Migration Service

- Build portable Ark appliances
- Temporarily installed to non-disruptively migrate data
- Could also be used for testing, performance measurement, certification
- Sold for customer or vendor use
- Professional Services

ARK SYSTEMS™

Supporting Information

Front End (Host Interfaces)

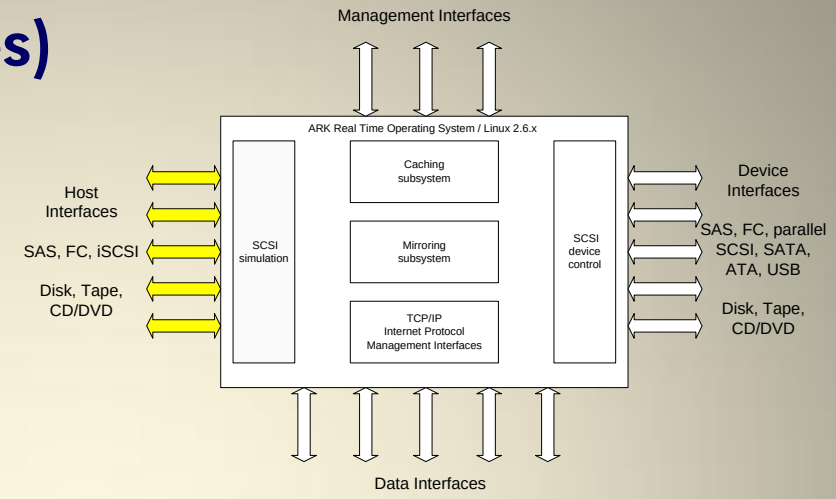

- Any number of interfaces / ports
- SAS, Fibre Channel, iSCSI, SAS, FCoE, parallel SCSI, and SATA ports
- Port types may be intermixed – volumes can be shared amongst protocols
- Fibre loop, point-to-point, fabric supported, fabric name server supported
- SAS expanders supported
- Multiple port iSCSI supported, iSNS supported

- A logical unit may be shared on any single or group of interfaces
- Different logical units may be accessed by different host(s) on the same port
- Logical unit characteristics may differ based on port / host definitions

- Ports may also act as initiators, allowing for extended copy functions

ARK SYSTEMS™

SCSI Device Emulation

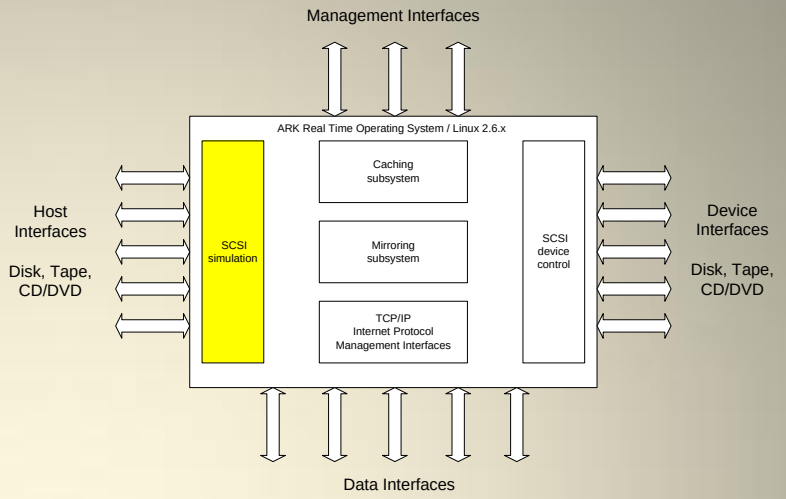

- SCSI emulation is independent of physical port characteristics
- Multiple device types supported (disk, tape, CD, SES)
- Most complete emulation, ever
- Rigorously tested using multiple conformance tests
- Supports “all” hosts

- Fast path for cache “hits”
- Supports “pass-through” mode for physical device management

- 16-32 byte data CDB’s and 64-bit block addresses
- Protection information / DIF support

ARK SYSTEMS™

Caching Subsystem

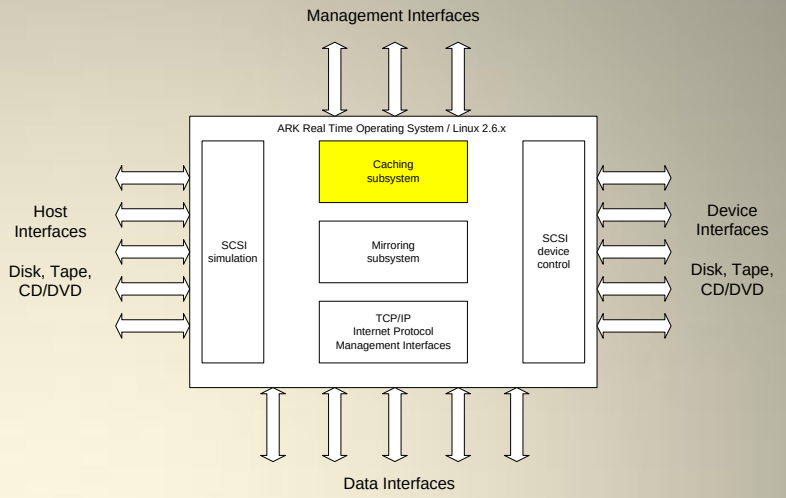

- "Store-in" cache – may be bypassed
- 32 and/or 64-bit cache memory addressing
- Blocks located via adaptive hashing tables

- "No single point of failure" via cache shadowing to other nodes
- 4KB "sparse" cache blocks
- Preemptive pre-fetch based on sequential activity
- Preemptive pre-fetch based on filesystem

- Nonvolatile "solid state disk" supported
- Collating, "lazy writes" supported
- Block range locking
- FLASH storage / devices

ARK SYSTEMS™

Virtualization and Mirroring

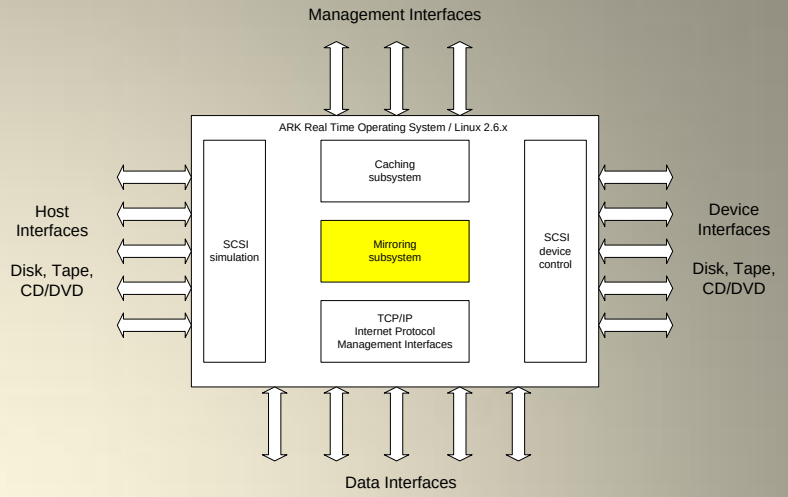

- Volume physical data may be located on any node
- Data may span multiple physical volumes – block granularity
- "N" volume mirrors – async, sync, semi-sync
- Nonvolatile change maps maintained to minimize re-establish time
- Automatic suspend/reestablish for performance
- Cascaded mirrors
- Volume consistency groups
- Facility, device, block recovery from mirrored volumes
- "N-level" snapshots
- "Virtual shared volumes"
- Many, many tuning knobs
- Encryption / RAID implemented via hardware

ARK SYSTEMS™

Back End

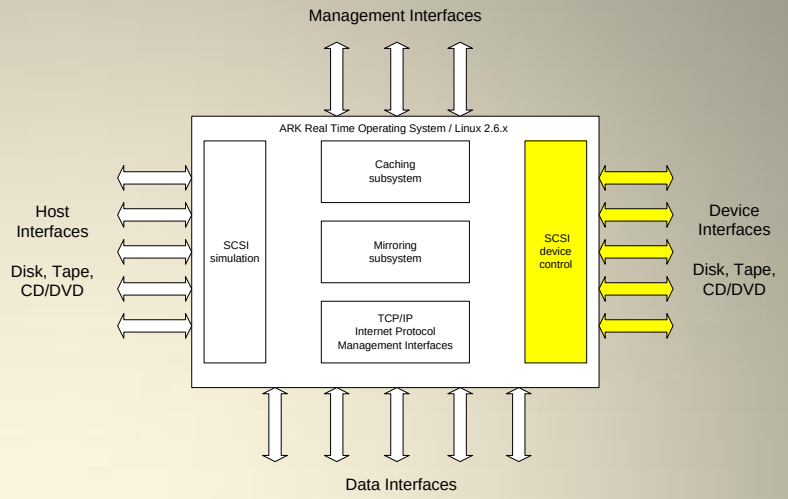

- SAS, Fibre Channel, SATA, parallel SCSI, iSCSI, USB, FLASH
- Caching subsystem makes SCSI requests, back end translates
- Many vendor-unique functions supported
- Read requests may fetch data from any mirror
- Multiple paths to devices – load balancing, recovery
- Multiple device types supported
- Built-in, online functionality, performance, and reliability testing
- Device scrubbing

ARK SYSTEMS™

Data Side

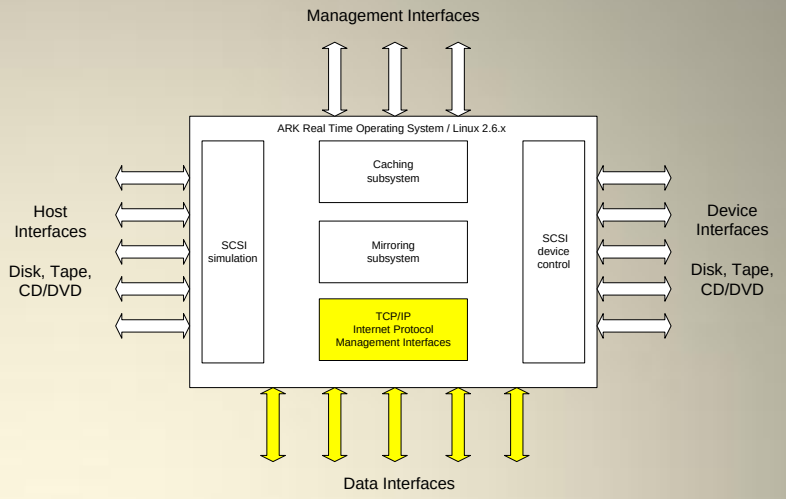

- Ports used for communications amongst Ark nodes
- Transfers mirrored data, coherent cache data, metadata
- Fibre Channel for “campus” distances, TCP/IP for “any” distance
- Custom TCP/IP stack optimized for low latency block data transfers
- TCP/IP stack supports and load balances multiple, simultaneous links
- FC ports can also be used for host/device access
- TCP/IP ports can also be used for management
- IPv6 in process

ARK SYSTEMS™

Management Interface

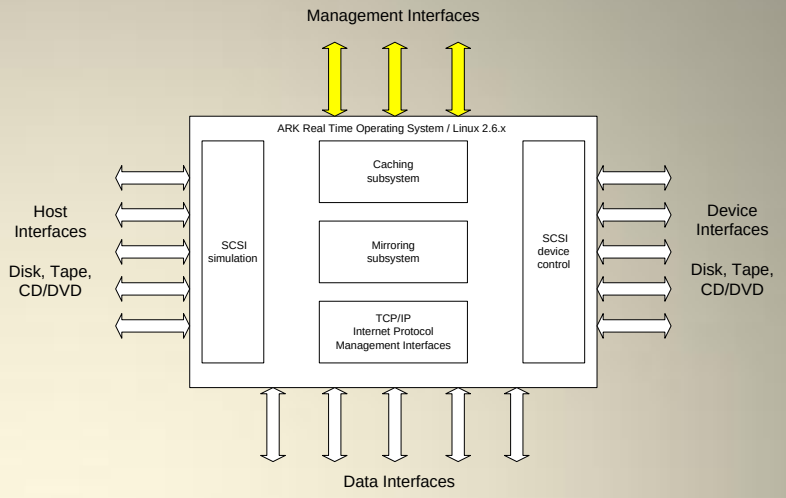

- Management is “out of band”
- Supports “standard” TCP/IP protocols
- Firmware, logs, “files” transferred via FTP
- Command/control via Telnet
- HTTP server
- SNMP monitoring, alerting, and control
- Environmental monitoring, supports IPMI
- TCP/IP authentication via source IP, user, password
- Four levels of command, file transfer privilege

ARK SYSTEMS™

Operating System

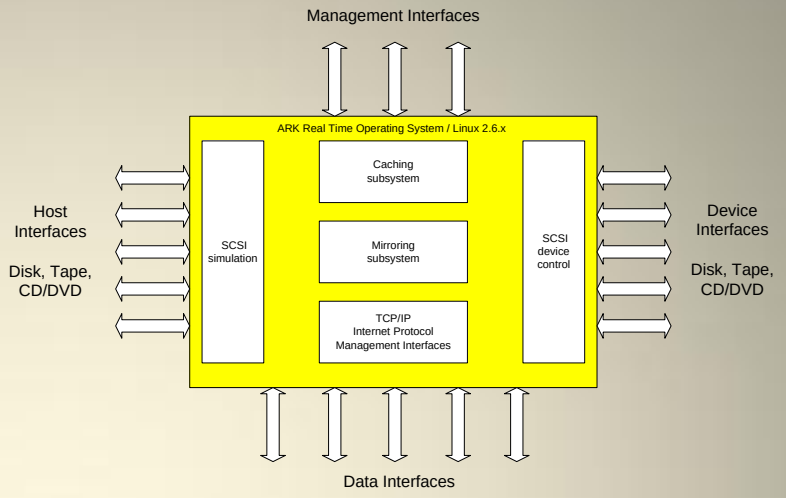

- Real time, multiprocessing, multitasking operating system
- Written to be portable
- Written to be platform agnostic
- Average time to port to a new x86 platform – one day
- "Zero-based budgeting" architecture – only what was needed was implemented
- Build process optimized for quick development – code change to test < 30 secs.
- Copious debugging, tracing, diagnostic, performance monitoring facilities

ARK SYSTEMS™

The "Appliance"

- Advanced storage functionality for low-end customers
- Built from "off the shelf" hardware and a storage director application
- Together, these components create in-band storage appliances, or directors
- These products add advanced functionality such as local and remote data access and mirroring, data migration, and data mining to existing heterogeneous servers and storage
- Customers easily add Ark products to existing installations, as neither drivers nor configuration changes are required
- Market now accepts "in-band" appliances

ARK SYSTEMS™

Lines of Code (RTOS + App)

Source: LocMetrics

Symbol	Count	Definition
Source Files	393	Source Files
Directories	18	Directories
LOC	609,823	Lines of Code
BLOC	52,900	Blank Lines of Code
SLOC-P	463,506	Physical Executable Lines of Code
SLOC-L	249,645	Logical Executable Lines of Code
MVG	55,903	McCabe VG Complexity
C&SLOC	145,295	Code and Comment Lines of Code
CLOC	93,417	Comment Only Lines of Code
CWORD	1,015,252	Commentary Words
HCLOC	7,505	Header Comment Lines of Code
HCWORD	38,892	Header Commentary Words

IP Licensing

- Filesystem-based Prefetch
- Virtual Shared Volumes
- License implementation-independent software
- Differentiating performance boost to licensee
- Or, license/sell everything to somebody
  - Storage vendor licensing outside IP
  - Start-up looking for IP to save time-to-market

Patents

1. 7,082,390 Advanced storage controller
2. 7,080,207 Data storage apparatus, system and method including a cache descriptor having a field defining data in a cache block
3. 6,931,566 System and method for optimizing fault tolerant storage controllers
4. 6,892,277 System and method for optimizing remote data content distribution
5. 6,862,671 System and method for optimizing establishment of mirrored data
6. 6,738,937 Method for nondisruptive testing of device and host attachment to storage subsystems
7. 6,499,091 System and method for synchronizing data mirrored by storage subsystems
8. 6,363,462 Storage controller providing automatic retention and deletion of synchronous back-up data
9. 6,360,306 Relocation of suspended data to a remote site in a distributed storage system
10. 6,345,368 Fault-tolerant access to storage arrays using active and quiescent storage controllers
11. 6,282,610 Storage controller providing store-and-forward mechanism in distributed data storage system
12. 6,073,209 Data storage controller providing multiple hosts with access to multiple storage subsystems

ARK SYSTEMS™