

# Storage Security – From Research to Industry Best Practice

Erik Riedel

Director, Interfaces & Architecture

Seagate Research

NASA Goddard, March 2008



# Seagate Research

- The Research Center currently has a total of 150 employees, all in Pittsburgh, Pennsylvania.
- About 120 technical staff from over 20 countries; almost 100 staff with Ph.D. degrees from over 50 universities.
- Built up from just two people in a rented office since 1999.



**Seagate Research**  
1251 Waterfront Place  
Pittsburgh, PA



# Outline

- Motivation
- Technology
- Standards
- Partnerships
- What's Next

# Motivation



terminology presented  
at FAST 2002, see  
survey paper in same  
conference, credit to  
my co-authors

a framework for  
evaluating storage  
system security

mahesh kallahalla,  
erik riedel,  
ram swaminathan

hp labs  
january 2002

# Motivation – Attacks on Stored Data

Attacks (* the fraction of system managers reporting the listed attack to the CSI/FBI survey in 2001 and 2006)	2001*	2006*	leak data	tampering	destroy data	revoked user
virus	94%	65%	-	-	X	-
laptop/mobile theft	64%	47%	X	-	X	-
insider abuse of net access	--	42%	-	-	-	-
unauthorized access to information	49%	32%	X	X	-	X
denial of service	36%	25%	-	-	-	-
system penetration	40%	15%	X	X	X	-
theft of proprietary information	26%	9%	X	-	-	X

# Motivation – Cost of Security

- “Users will not pay for [storage] security technology” – storage marketing circa 2001
- “Almost every customer asks about our security strategy” – storage vendor circa 2006
- Average annual spending on computer security is \$200 – \$1,350 per employee\*

\*CSI/FBI 2006 Computer Crime Security Survey

# Technology

\*incubated by Seagate Research (Thibadeau) in conjunction with multiple Seagate technology and product teams

# Technology – Basics

- Leak of data – encryption
- Tampering – hashing  
(turns it into destruction)
- Destruction – difficult to fix directly  
(make backups/replicas)
- Revoked users – careful key mgmt
- Denial of service – no great solutions  
(replicas are a start)

# Technology – Toolbox

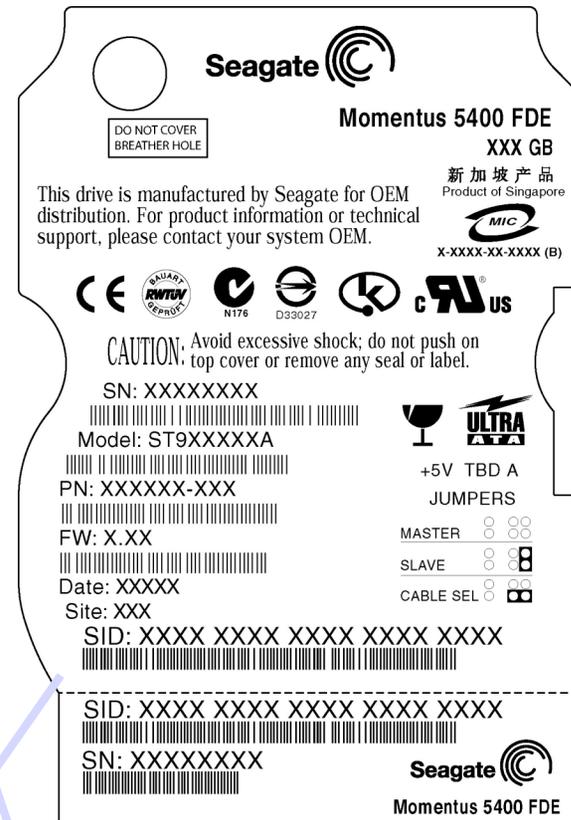
- Additional relevant features
    - Key generation via random #s
    - Key exchange over networks
    - Roots of trust & attestation
    - Audit logs & data provenance
- for key mgmt
- trusted computing base
- via segregated storage



These technology **mechanisms** must be tied together into secure end-to-end **solutions**

# Technology – DriveTrust

- Seagate FDE
  - Full disc encryption at native data speed
  - No performance penalty
  - Secure partitions w/ fine grain access control
  - Crypto functions
  - Root of trust
  - Command set on SCSI and ATA

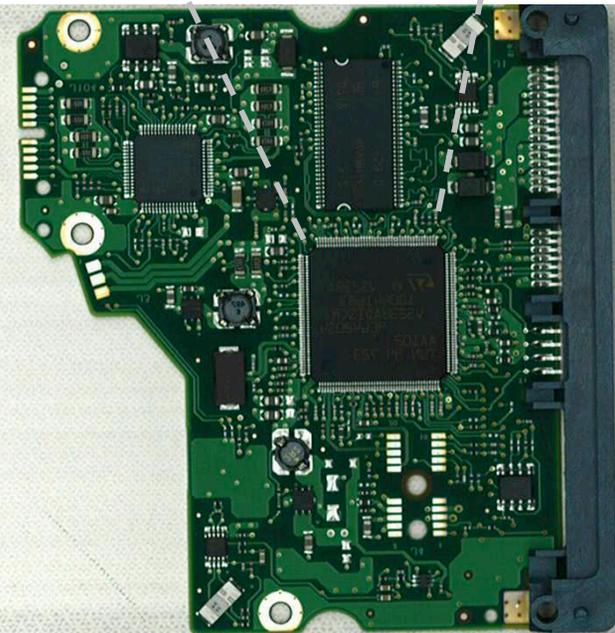


Security ID on permanent & perforated label

# Technology – Electronics



- high-speed, specialized data movement & protocol processing



# Technology – DriveTrust

- Data is stored encrypted
- Encrypted data never leaves the drive
- Repurposing done via fast secure delete
- Access control
  - Partition access managed by SPs
  - Pre-boot authentication
- Hidden areas managed via TCG commands
  - Personal data, License keys, other



# Standards

# Standards

- Trusted Computing Group (TCG)
  - Estimated 25% of 2006 PCs shipped with a TPM; predicted 50% of 2007 shipments
- Storage WG – includes 6 disk drive makers
  - Key use cases
    - Enrollment – device pairing via public/private keys
    - Encryption of user data
    - Private, hidden storage areas
  - All functions independent of the operating system, managed in a separate computing environment
- T10 (SCSI) and T13 (ATA) command sets

# Partners

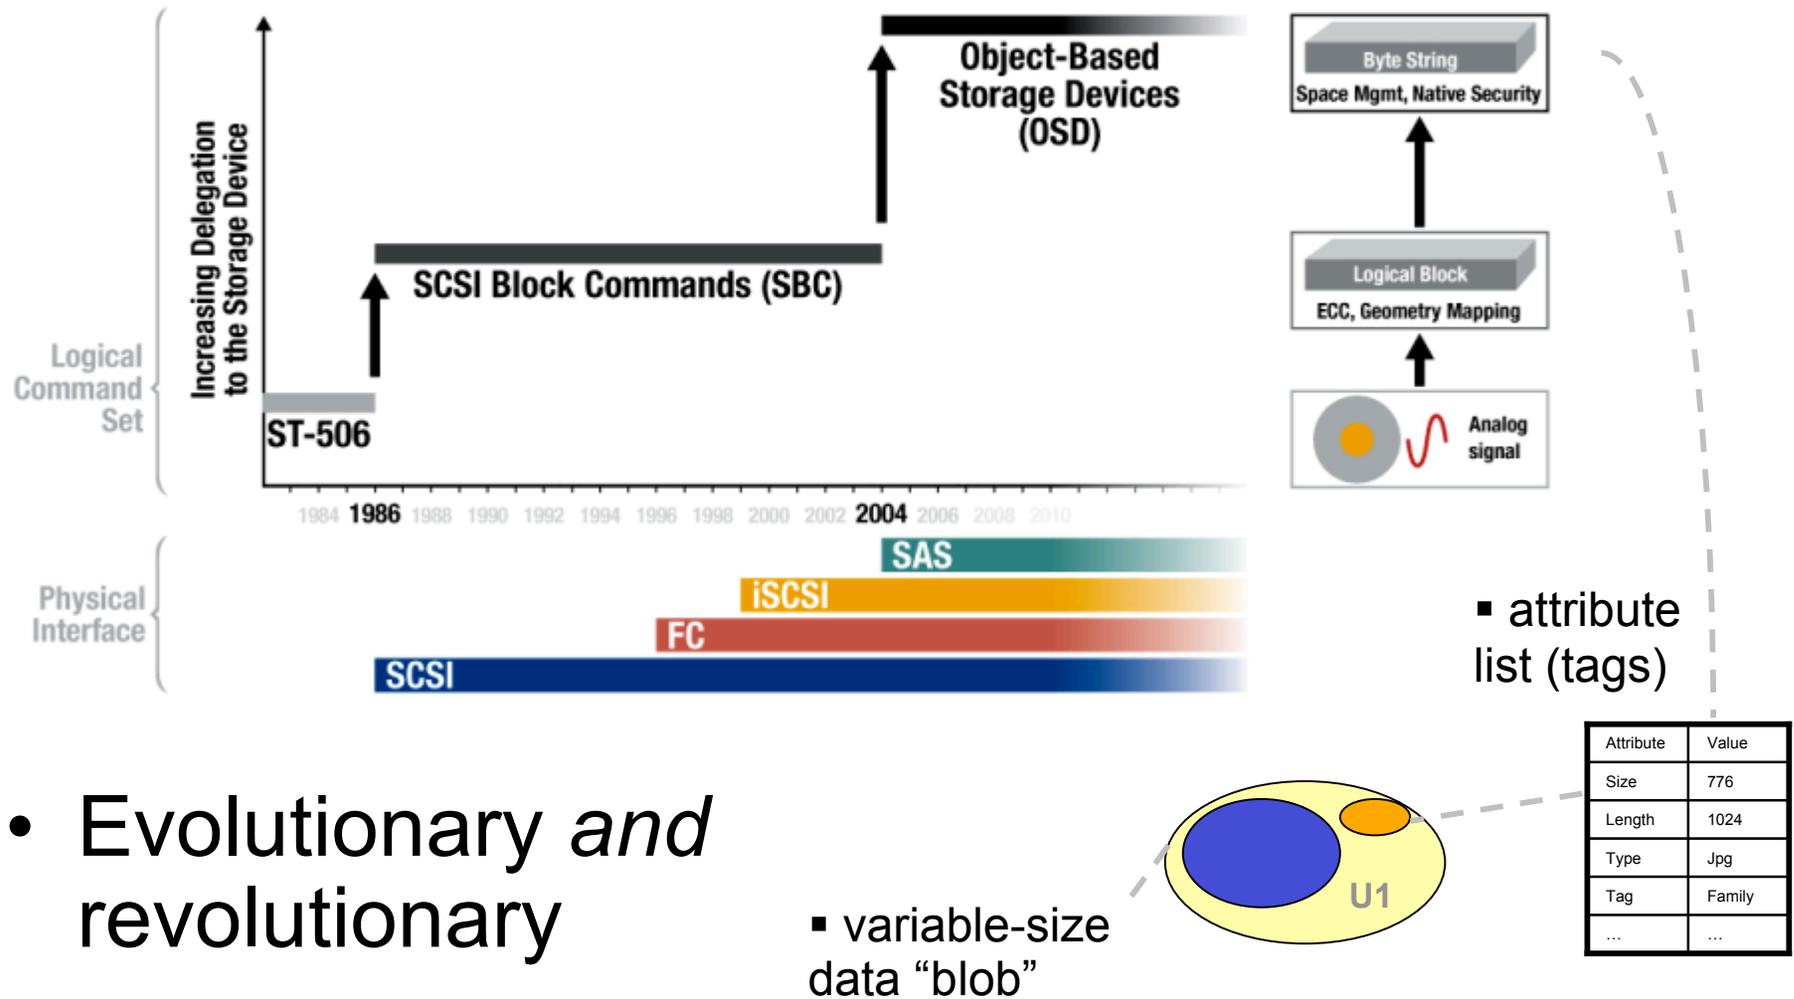
# Partners – DriveTrust

- Burned basic mechanisms into silicon
  - Full interface speed
  - Isolated from host
  - Integrated with data access
- Leave control and management to software APIs
- Enable a variety of software partners
  - Wave Systems (US); Secude (Switzerland); CryptoMill (Canada); GuardianEdge (US); seeking additional partners



# What's Next

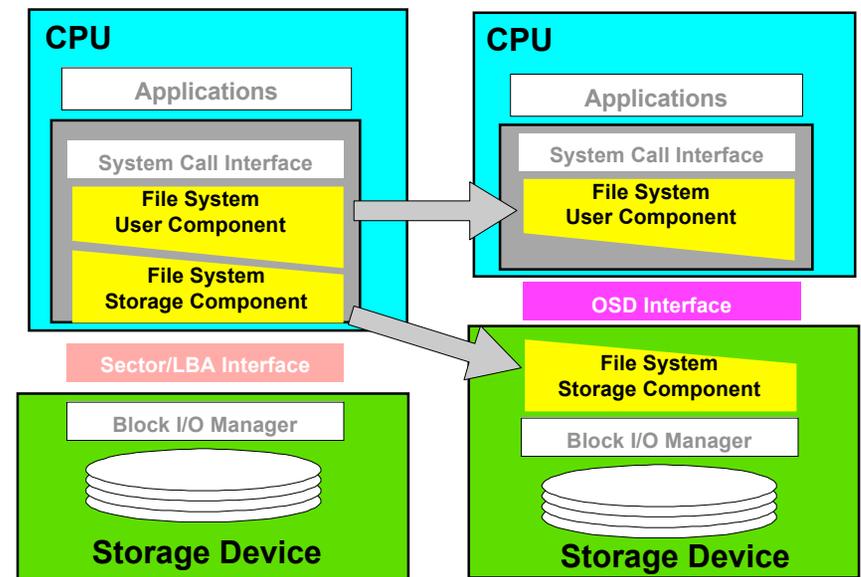
# What's Next – Objects



- Evolutionary *and* revolutionary

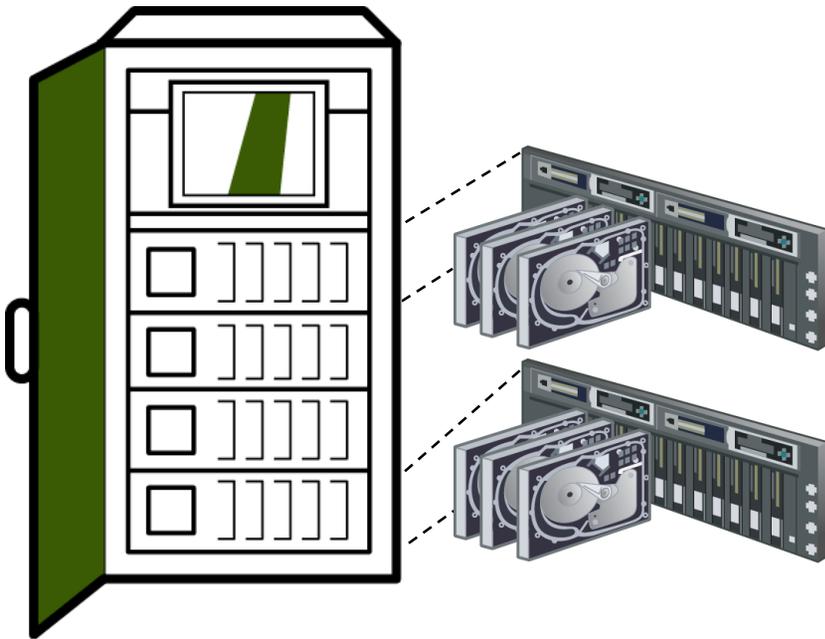
# What's Next – Objects (2)

- Object-based Storage Device (OSD) standard developed by SNIA and T10
  - OSD-1 (September 2004); OSD-2 (Early 2008)
- Space management handled by individual devices (arrays, controllers, drives); variable sized objects
- Fine-grained security via a shared key system
  - per-object capabilities
  - per-request authorization
  - authentication & access control at security manager
- Optional in-flight integrity
- At-rest protection possible
  - with hardware acceleration



# Enterprise Objects

- High-density and high-capacity arrays being built by all of the system vendors



- 15 drives in 3U (hot-swap)
  - over 200 drives per cabinet
  - 200 TB @ 1 TB/drive
- 112 drives in 3U (power-managed)
  - up to 900 drives per cabinet
  - almost 1 PB @ 1 TB/drive
- Used by large enterprises
  - goes by many names
  - clusters; grids; data mining
- Systems will benefit from delegated management
  - offload expensive mechanisms (encryption, compression, search, ...)
  - policy managed centrally

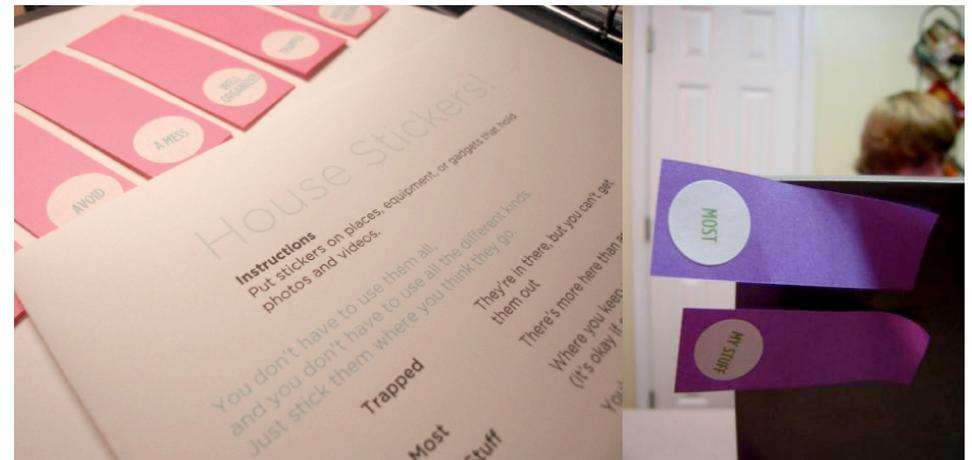


## For our Home Visits project: “Sticker tour”

We gave people sheets of stickers, and asked them to “tag” things related to photos, videos, and music with different labels. Then they gave us a tour of the stickers, **leading us to key objects and places and giving us stories associated with each.**

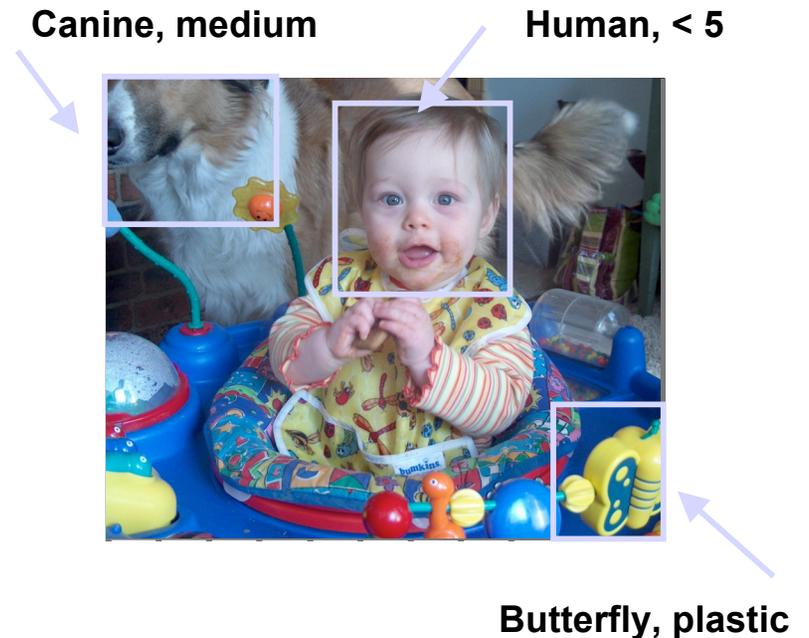
### The Stickers

Trapped  
My Stuff  
Most  
Private + Personal  
Avoid  
A Mess  
Well-organized



# What's Next – Data Aware

- Increase semantic understanding of the stored data
  - automated extraction of metadata tags
  - combine with existing user-created tags
  - apply image and video recognition algorithms
  - apply specialized electronics



# What's Next – Privacy

- Infroperty – information as property
  - Premise – privacy is gone for today's data, but there is hope for data created in the future
  - Legal status needs to catch up with the technology and the importance of digital data
  - Information and property merge
  - Personal information bound to an infroperty agent, agreements for use negotiated by agents wherever data is used
- Must include a relevant legal framework
  - See [www.istpa.org](http://www.istpa.org)
  - Int'l Security Trust & Privacy Alliance

# Conclusion

- We've made a start, but have a long way to go
- Advice to researchers – find relevant user problems, consider paths to solution, stick with it
- Advice to companies – work more closely with researchers for novel solutions, novel views on the problem
- Looking for new software functions and partners to use our mechanisms

**Seagate**



# Backup Slides

# Technology – Trade-Offs



Barracuda 7200  
up to 1 TB

Over 40 unique  
drive models

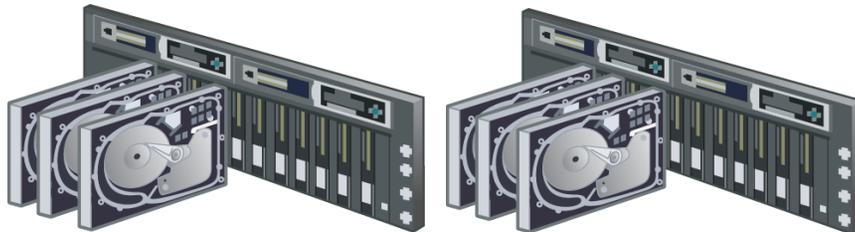


Cheetah 15K  
up to 300 GB

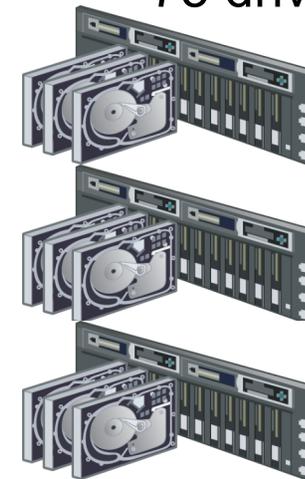
Savvio 10K  
up to 146 GB



30 drives / 6U



75 drives / 6U



# Technology – Design Points



**DB35**  
high-cap, quiet



**EE25**  
small, shakable



**ST18**  
smaller, lower power



**Momentus FDE**  
small, secure



**Momentus PSD Hybrid**  
high-perf, low power



**Pocket,  
Portable  
carryable**



framework

players

- *owners*
  - create data
  - determine access to data
- *readers* – read
- *writers* – modify
- *storage servers*
  - store/retrieve bits
- *group servers* (many flavors)
  - handle “delegated” keys
- *adversaries*
  - tampers with data
  - may collude w/ others

# threats and attacks

attacks, as reported in survey of system managers by CSI/FBI, Spring 2001 <small>*of ~500 responses, 78% had financial losses, only 37% could estimate damage</small>	% surveyed	damage (\$ millions)*	msgs		data			revoked user	denial of service
			leak	change	leak	change	destroy		
telecom eavesdropping	10%	1	✓	::	::	::	::	::	::
active wiretap	2%	n/m	-	✓	::	::	::	::	::
system penetration	40%	19	✓	✓	✓	✓	✓	::	::
laptop theft	64%	9	-	::	✓	::	✓	::	::
theft of proprietary info	26%	150	-	::	✓	::	::	✓	::
unauth access by insiders	49%	6	-	::	✓	✓	::	✓	::
sabotage	18%	5	-	::	::	::	✓	::	✓
virus	94%	45	-	::	::	::	✓	::	::
denial of service	36%	4	-	::	::	::	::	::	✓

framework

attacks

- attacks on data
  - leak
  - change
  - destroy
- adversary
  - act alone
  - collude w/ server
  - revoked user
- compromise group server
- denial of service

# security guarantees - existing systems

system	message attacks	adversary			w/ storage srv			revoked		subvert group server	denial of service
		leak	change	destroy	leak	change	destroy	leak	change		
CFS	-	✓	✓	X	✓	✓	X	-	-	-	X
SFS-RO	✓	✓	✓	X	✓	✓	X	X	-	✓	X
Cepheus	✓	✓	✓	✓	✓	✓	X	✓	✓	X	X
SNAD	✓	✓	✓	✓	✓	✓	X	✓	✓	X	X
NASD	✓	✓	✓	✓	X	X	X	✓	✓	X	X
iSCSI w/ IPsec	✓	✓	X	X	X	X	X	✓	✓	-	X
LUN security	X	X	X	X	X	X	X	X	X	-	X
AFS	✓	✓	✓	✓	X	X	X	✓	✓	✓	X
NFSv4	✓	✓	✓	✓	X	X	X	✓	✓	X	X
PASIS/S4	-	-	-	✓	✓	✓	✓	-	-	-	X
OceanStore	-	✓	✓	✓	✓	X	✓	✓	✓	-	X