Data Integrity

Means and Practices

Raymond A. Clarke
Sr. Enterprise Storage Solutions Specialist, Sun Microsystems - Archive & Backup Solutions
SNIA Data Management Forum, Board of Directors

Backup vs. Archiving – there's a difference

Both are required in today's environments

BACKUP

Sinngle/Multiple copies
Multiple points in time

Recover data/information
Due to corruption or loss
Meet RPO and RTO objectives
Maintain copy for disaster recovery
Offline volume remounted and manually searched

SSD
Replication
High Performance Disk
Encryption
Capacity Disk
De-duplication
VTL - ATL

Primary Data

ARCHIVING

Multiple copy
Infinite time periods

Maximize efficiency and optimization
Regulatory compliance, Provenance, Fixity
Enable eDiscovery
Meet best practice
Search Criteria Online files recalled based on key word/date criteria

Disk Archive
Tape Archive

Replication
Encryption
VTL – ATL
Deep Archive

Why is Backup & Archive So Important?

... because The History of Data Growth is Exponential!

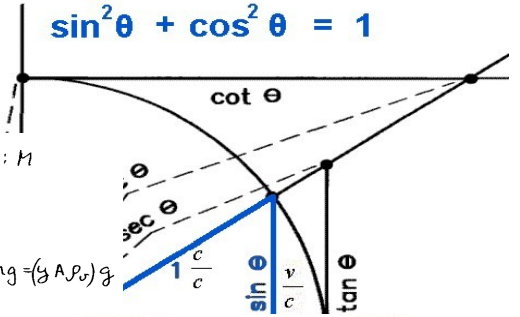

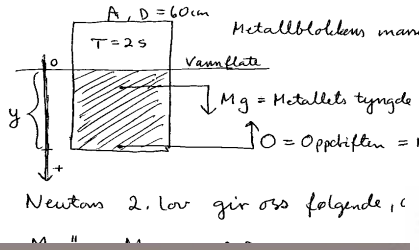

24 Words - Pythagorean Theorem

67 Words - Archimedes Principal

179 Words - 10 Commandments

286 Words - Lincoln's Gettysburg Address

1300 Words - US Declaration of independence

26911 Words .....
EU REGULATION ON THE SALE OF CABBAGES

Building a Terminology Bridge

Archive: the report advocates that IT practices adopt a more consistent usage of the term 'archive' with other departments within the organization. To the archival, preservation, and records management communities, **an "archive" is a specialized repository with *preservation services and attributes*.**

Preservation: managing information in today's datacenter with requirements to safeguard information assets for eDiscovery, litigation evidence, security, and regulatory compliance requires that many classes of information be preserved from time of creation. ***Preservation is a set of services that protect, provide availability, integrity and authenticity controls, include security and confidentiality safeguards, and include an audit log, control of metadata, and other practices for each preservation object*.** The old IT practice of placing information into an archive when it becomes inactive or expired no longer works for compliance or litigation support, and only adds cost.

Authenticity: is defined in a digital retention and preservation context as a practice of verifying a digital object has not changed. ***Authenticity attempts to identify that an object is currently the same genuine object that it was "originally" and verify that it has not changed over time unless that change is known and authorized***. Authenticity verification requires the use of metadata. The critical change for IT practices is that metadata is now very important and must be safeguarded with the same priorities the data is. IT practices

http://www.snia.org/forums/dmf/knowledge/term_bridge/

Source: SNIA Data Management Forum

What is an Archive?

A Searchable Repository That Provides Business Benefits

- Security
- Accessibility
- Integrity
- Scale
- Long Life
- Open Standards (Access and data format)
- Cost and “Data” Effective
- Eco Responsible

Demands of a New Archive Reality

Is the ratio for archiving solutions changing?

10 / 90

versus

2 / 18 / 80

- Next Generation Archives need to address a new dimension of the massive resting data – How do you search Petabytes of data from the edge?
- The new ratio has evolved into a Write / Read / Search relationship (2 / 18 / 80) – *different demands on the infrastructure*
- Business semantics need to drive data management not systematic schemas
- Virtualization and Search become critical to the presentation of the data, something new is needed...
- Compute and Store need to Converge

Most Data Remains Untouched

Average Distribution of Data by Tier

Age in Days	Probability of Re-reference
1	70-80%
3	40-60%
7	20-25%
30	1-5%
90+	Near 0%

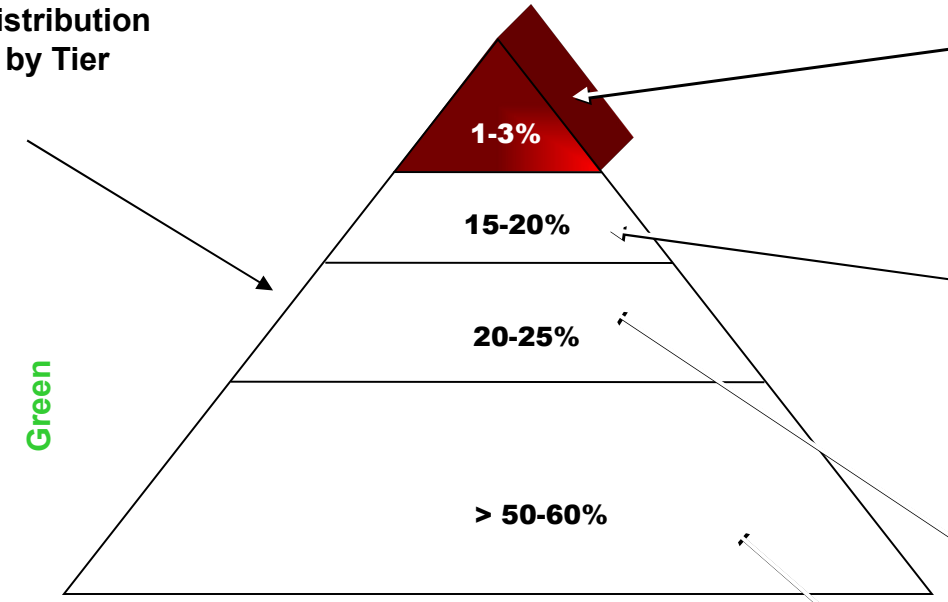

Tier 0
Ultra High-performance/Ultra High value Information

Tier 1
High-value, High Ingest, OLTP, Revenue Generating, High-performance Data

Tier 2
Backup/recovery Apps, Reference data, Vital and Sensitive Data, Lower value active data

Tier 3
Fixed Content, Compliance, Archive, Long-term Retention, Green Storage Apps

Value Index %	Type of Technology
T0 – 99.999+	DRAM SSD, Flash Memory HDD, Hi-Perf Disk
T1 - 99.999+	Enterprise-class HDD, RAID, Mirrors, Replication
T2 – 99.99	Midrange HDD, SATA, Virtual Tape, MAID, Integrated Virtual Tape Libraries
T3 – 99.9	High- Capacity Tape, MAID, Manual Tape, Shelf Storage

Source: Horison Information Strategies Fred Moore www.horison.com

Why Tape Continues to Make Good Sense

Function	Tape	Disk
Long span of media	15~30 years on all new media.	3~5 years for most HDDs
Portability	Media is completely removable and easily transported.	Disks are difficult to remove and safely transport.
Move data to remote location for DR with or without electricity	Data/Media can be move remotely with or without electricity.	Difficult to move disk data to remote location for DR without electricity.
Inactive data does not consume energy	Green storage	Very rarely, except with MAID (questionable ROI).
Encryption for highest security level	Encryption available on essentially all tape drives types.	Available on selected disk products.

Make a Fool-Proof System and Nature comes up with a more creative Fool!

- Human Error is the most likely and unpredictable source of problems
- The smartest people sometimes are the most likely to make an error
- How a well-designed system provides mitigation
  - Consider and mitigate all possible failure scenarios
  - Provide user-friendly, simple management interface
  - Eliminate human interaction as far as possible by policy-driven automated processes
  - Use Quorum to validate critical actions

"the smarter the person, the dumber the mistake"

Store Data for Forever!

Future-proof Data Storage for data preservation

- Archive files are self-describing, standard
- No lock-in, open TAR format
- Move data to newer, more reliable media over time *transparently*
- WORM enforcement throughout the archive

System Basics

- User/Application Storage Layer Abstraction
  - New Data
  - Aged Data
- Policies
- Multi-Tiered, Multi-copy Archival
  - Local
  - Remote
  - Distributed
  - Cacaded
- Continuous Data Protection, On-disk WORM and ***Encryption***.

Tape Encryption Technology

- Encryption Engine located between the Compression and formatting Functions
  - Encrypted Data is highly randomized so encryption must be done post-compression to retain the benefits of Compression
- All tape-based encryption products use AES-256 – the most powerful commercially available encryption algorithm
- All Firmware and Hardware encryption processes are validated by Known Answer Test at power-on
- Drive is designed to ensure that data cannot be encrypted with a corrupted key

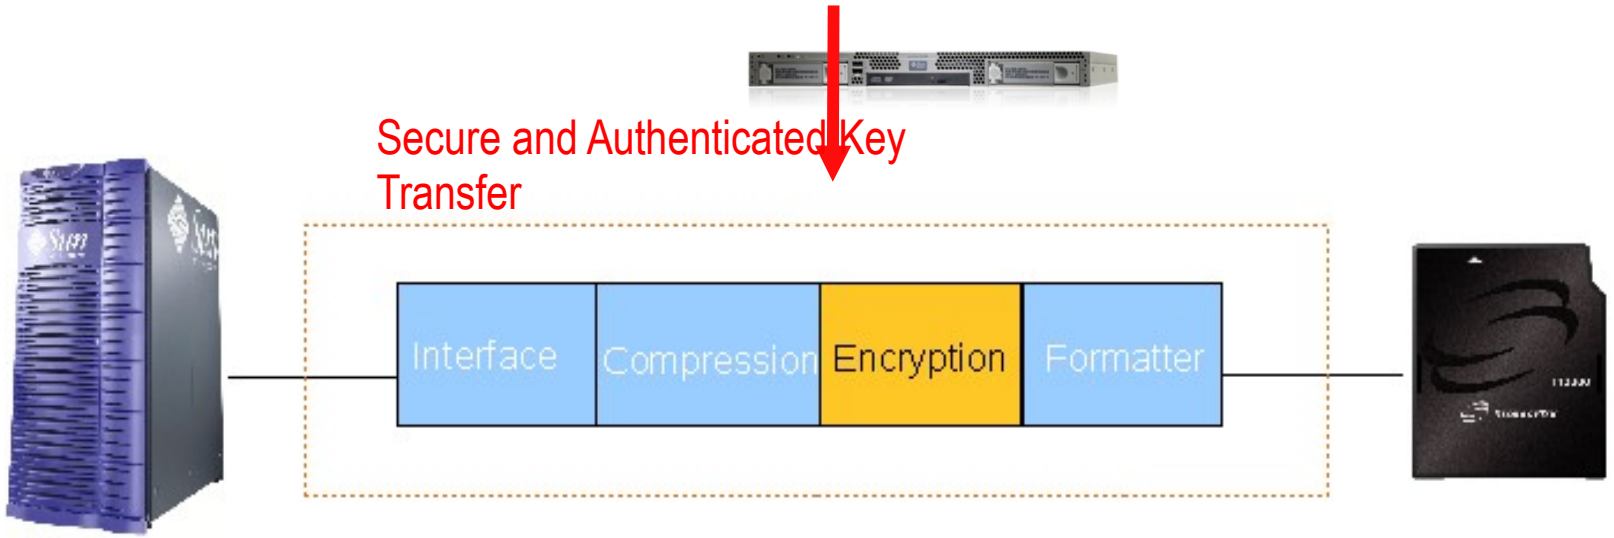

Typical Small Configuration

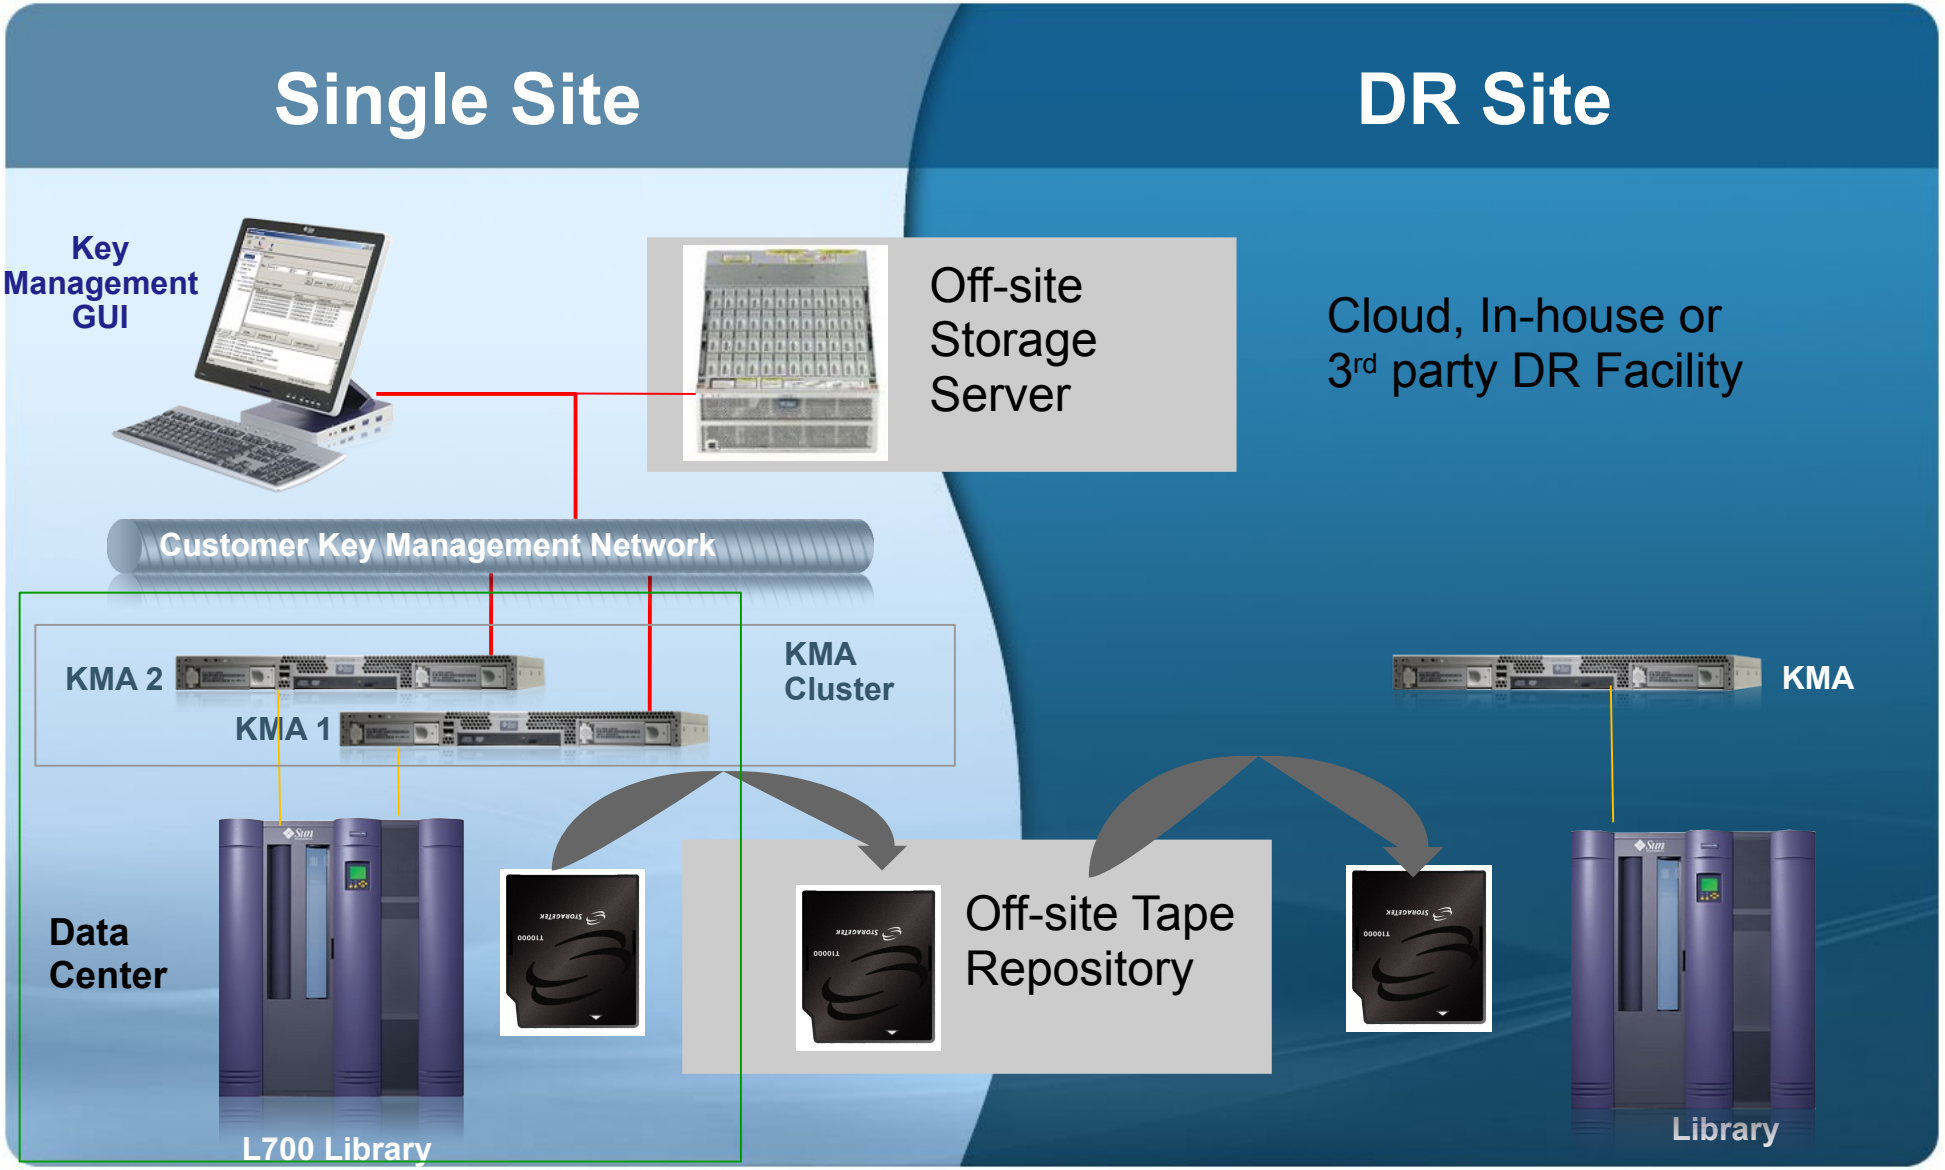

Key Life Cycle

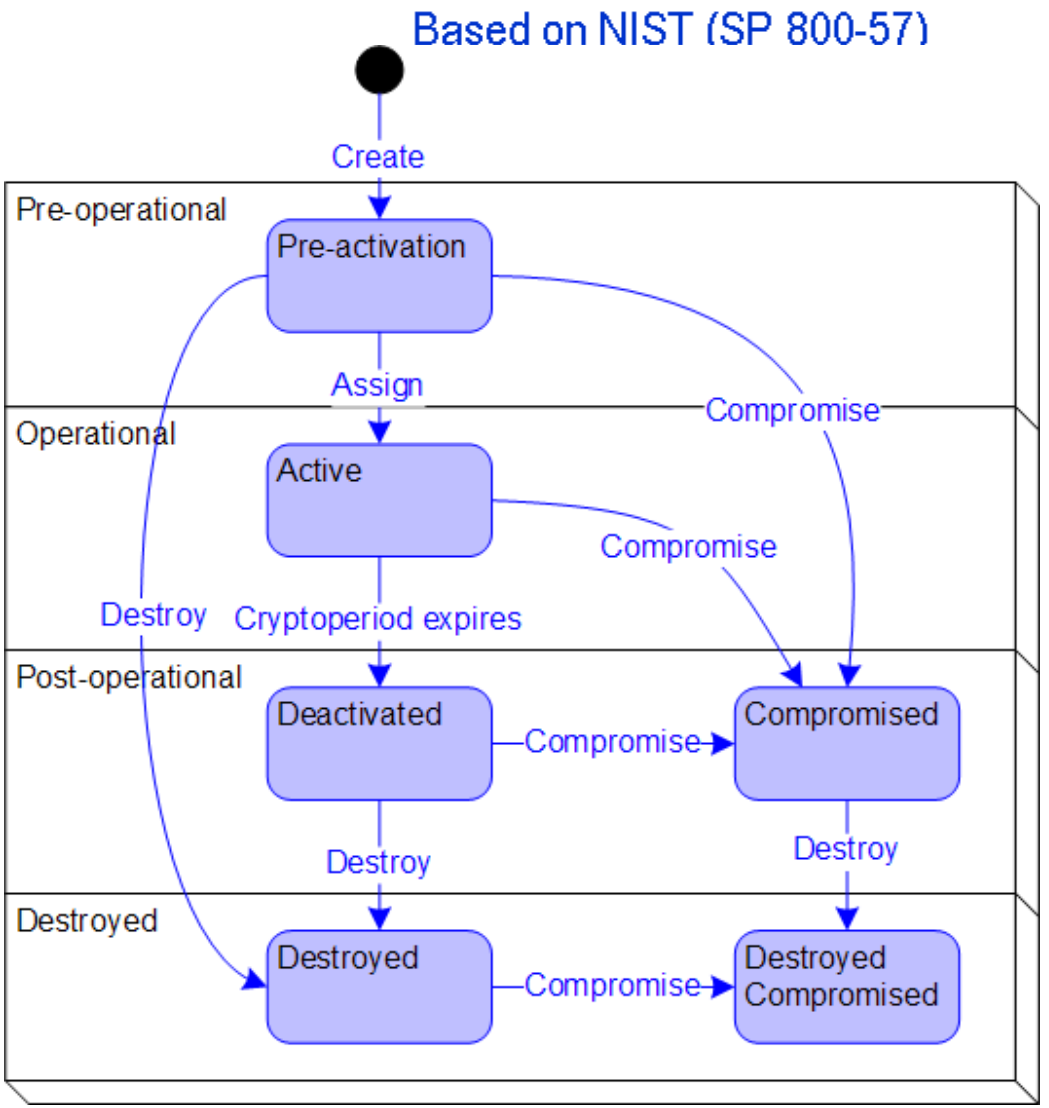

Each KMA maintains a reservoir of pre-activated keys that are replicated across the cluster

Key Life Cycle controlled by customer-defined Policy

Keys can be manually de-activated using Compromise function

Only keys in Post-operational State can be destroyed

What do we need to protect against?

Threat	Mitigation
Key Management Appliance Failure	KMS design replicates database to all KMA's in cluster. Database Backup protects universal multiple failures
Network Failure	KMS design can ride through temporary interruptions, managed switches can provide redundant network connection.
Data Center Fire, Flood etc.	KMS replication to off-site KMA's in cluster. Backup database to off-site server. Off-site tape vaulting.3rd Party DR Services.

Mitigation for Small Configuration

Single Site

Off-site Storage Server

Customer Key Management Network

KMA 2

KMA Cluster

L700 Library

DR Site

Cloud, In-house or 3rd party DR Facility

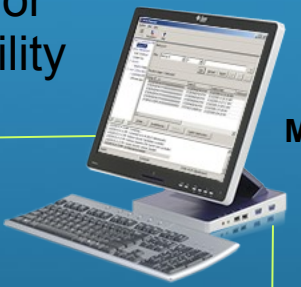

Key Management GUI

Database Transfer using Restore function (Quorum)

KMA

Off-site Tape Repository

Library

AES-256

- The most powerful commercially available algorithm
- AES-256 uses a 256-bit key
- A 256 bit number has $1.16X10^{77}$ permutations
- In July 2007, the population of the world was 6,602,224,175
- If you gave everyone in the world a super-computer that tries a key value every nanosecond, it would take $5.56 X 10^{50}$ years to try all combinations
- Assumes that key values are adequately random
- "At 20 to 30 x 10^9 years, the sun will expand into a red ball and die, overwhelming Earth with the heat. Oceans will boil and evaporate, and other planets near the sun also will burn" January 15, 1997
- With AES-256, it is imperative that your system protects itself against malicious or inadvertent loss of keys

FIPS 140-2 Security Levels

Modules are evaluated against 12 sets of criteria and assigned a Security Level

The Security Level of the Complete Module is determined by the lowest Security Level per criterion

- Security Level 1 is “Basic"
- Security Level 2 adds "Tamper Evidence" often by using approved labels.
- Security Level 3 is "Tamper Resistant" often by encapsulating the device in thick epoxy
- Security Level 4 is “Tamper Respondent" for example active circuitry will erase keys if anyone tampers with the device.

Sun T10000B FIPS Certificate

FIPS 140-2 provides four increasing, qualitative levels of security: Level 1, Level 2, Lev
the wide range and potential applications and environments in which cryptographic mo
cover eleven areas related to the secure design and implementation of a cryptographic
cryptographic modules as tested in the product identified as:

Sun StorageTek™ T10000B Encrypting Tape Drive by
(Hardware Version: P/N 315488302; Firmware Versions: 1.40.20

Prior to Certification of the Module, the implementation of each cryptographic algorithm used in the module must be tested and FIPS-certified

and tested by the Cryptographic Module Testing accredited laboratory: ***InfoGard Laborato***
is as follows: ***CRYPTIK Version 7.0***

Cryptographic Module Specification:	Level 2	*Cryptographic Module Ports and Interfaces:*	Level 2
Roles, Services, and Authentication:	Level 2	*Finite State Model:*	Level 2
Physical Security: (Multi-Chip Standalone)	Level 2	*Cryptographic Key Management:*	Level 2
EMI/EMC:	Level 2	*Self-Tests:*	Level 2
Design Assurance:	Level 2	*Mitigation of Other Attacks:*	Level N/A
Operational Environment:	Level N/A	*tested in the following configuration(s):*	N/A

The following FIPS approved Cryptographic Algorithms are used: **AES (Certs. #495, #647, #941, #942 and #967); DRBG (Cert. #6); HMAC (Certs. #398 and #540); SHS (Certs. #736 and #937); RSA (Cert. #334)**

The cryptographic module also contains the following non-FIPS approved algorithms: **AES (Cert. #941, key wrapping; key establishment methodology provides 256 bits of encryption strength); RSA (key wrapping; key establishment methodology provides 112 bits of encryption strength); MD5**

Overall Level Achieved: 2

Signed on behalf of the Government of the United States

Signature: DFD

Dated: July 16, 2009

Chief, Computer Security Division
National Institute of Standards and Technology

Signed on behalf of the Government of Canada

Signature:

Dated: JULY 13, 2009

Director, Industry Program Group
Communications Security Establishment Canada


Sun
microsystems
Thank You
for
Your Time and
Attention
Raymond.Clarke@Sun.com
(212) 558-9321