

Secure Archival is Hard...

Really Hard

Christopher Smith, Maliha Tabassum, Soumya Chowdary Daruru,
Gaurav Kulhare, Arvin Wang, Ethan L. Miller*[†], Erez Zadok

Stony Brook University, PureStorage, UC Santa Cruz[†]*

The Problem at a Glance

How to keep long-term archival data **secure**?

Confidentiality

Integrity

Availability

$O(100\text{yrs})$



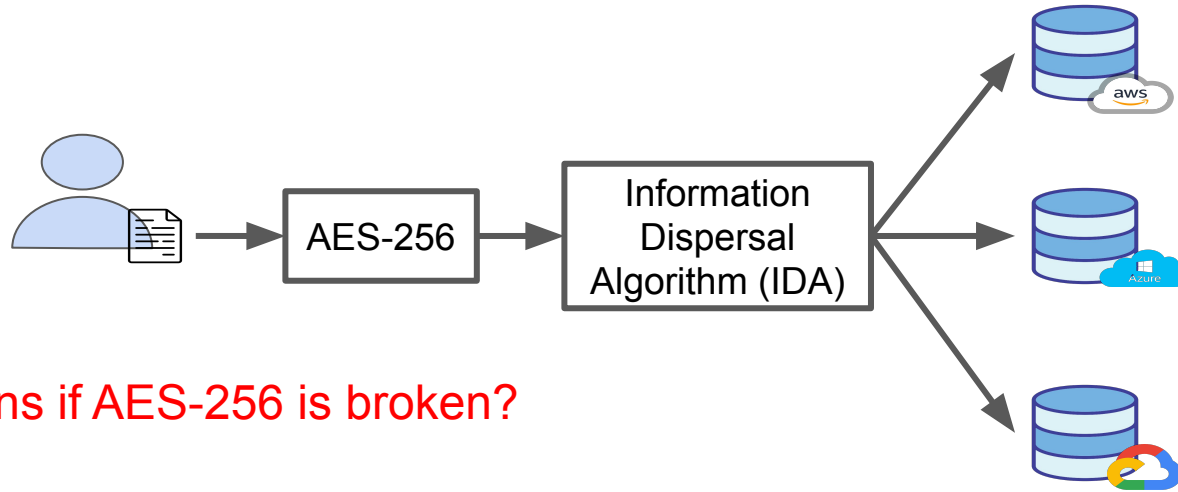
Use Cases

Challenges

- ❖ Cryptographic Obsolescence
- ❖ Harvest Now, Decrypt Later
- ❖ Storage cost
- ❖ Side-channel attacks
- ❖ And more...



Warm-Up Attempt



What happens if AES-256 is broken?

└─ Just re-encrypt...

└─ What if encrypted data is already stolen? // i.e. Harvest Now, Decrypt Later

Harvest Now, Decrypt Later

HARDWARE > QUANTUM | October 30, 2023

Are harvest now, decrypt later cyberattacks actually happening?

Cybercriminals may already be hoarding data for when quantum computers become powerful enough to break current encryption standards.

By Greg Noone

A REUTERS SPECIAL REPORT

U.S. and China race to shield secrets from quantum computers

Harvest now, decrypt later

 Add languages

Article [Talk](#)

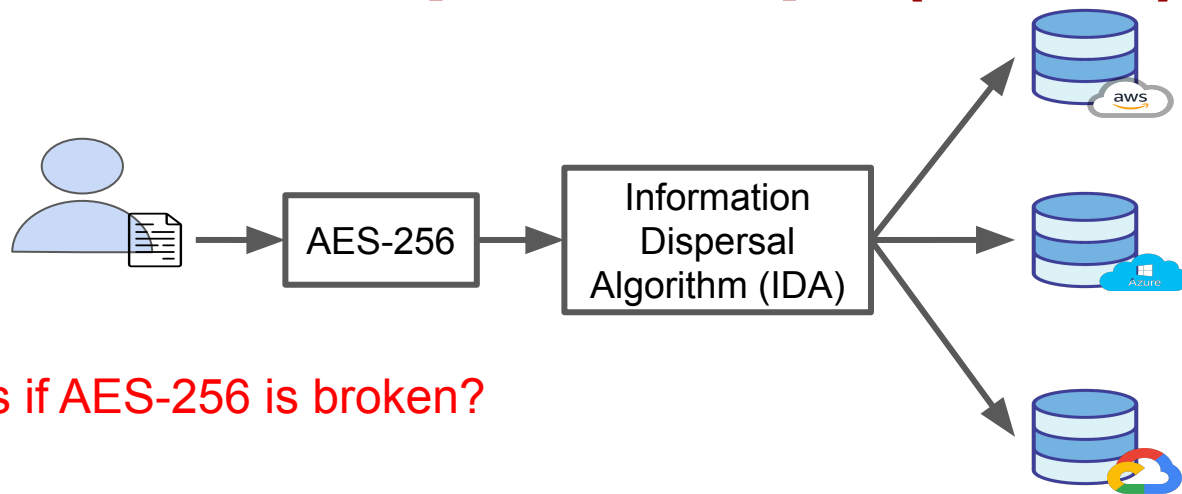
[Tools](#)

From Wikipedia, the free encyclopedia

Harvest now, decrypt later, also known as **store now, decrypt later** or **retrospective decryption**, is a surveillance strategy that relies on the acquisition and long-term storage of currently unreadable encrypted data awaiting possible breakthroughs in [decryption](#) technology that would render it readable in the future.^{[1][2]}

Harvest Now, Decrypt Later (HNDL) attacks are happening **now**

Warm-Up Attempt (cont.)



What happens if AES-256 is broken?

└ Just re-encrypt...

└ What if encrypted data is already stolen? // i.e. Harvest Now, Decrypt Later

└ Use an unbreakable encryption scheme???

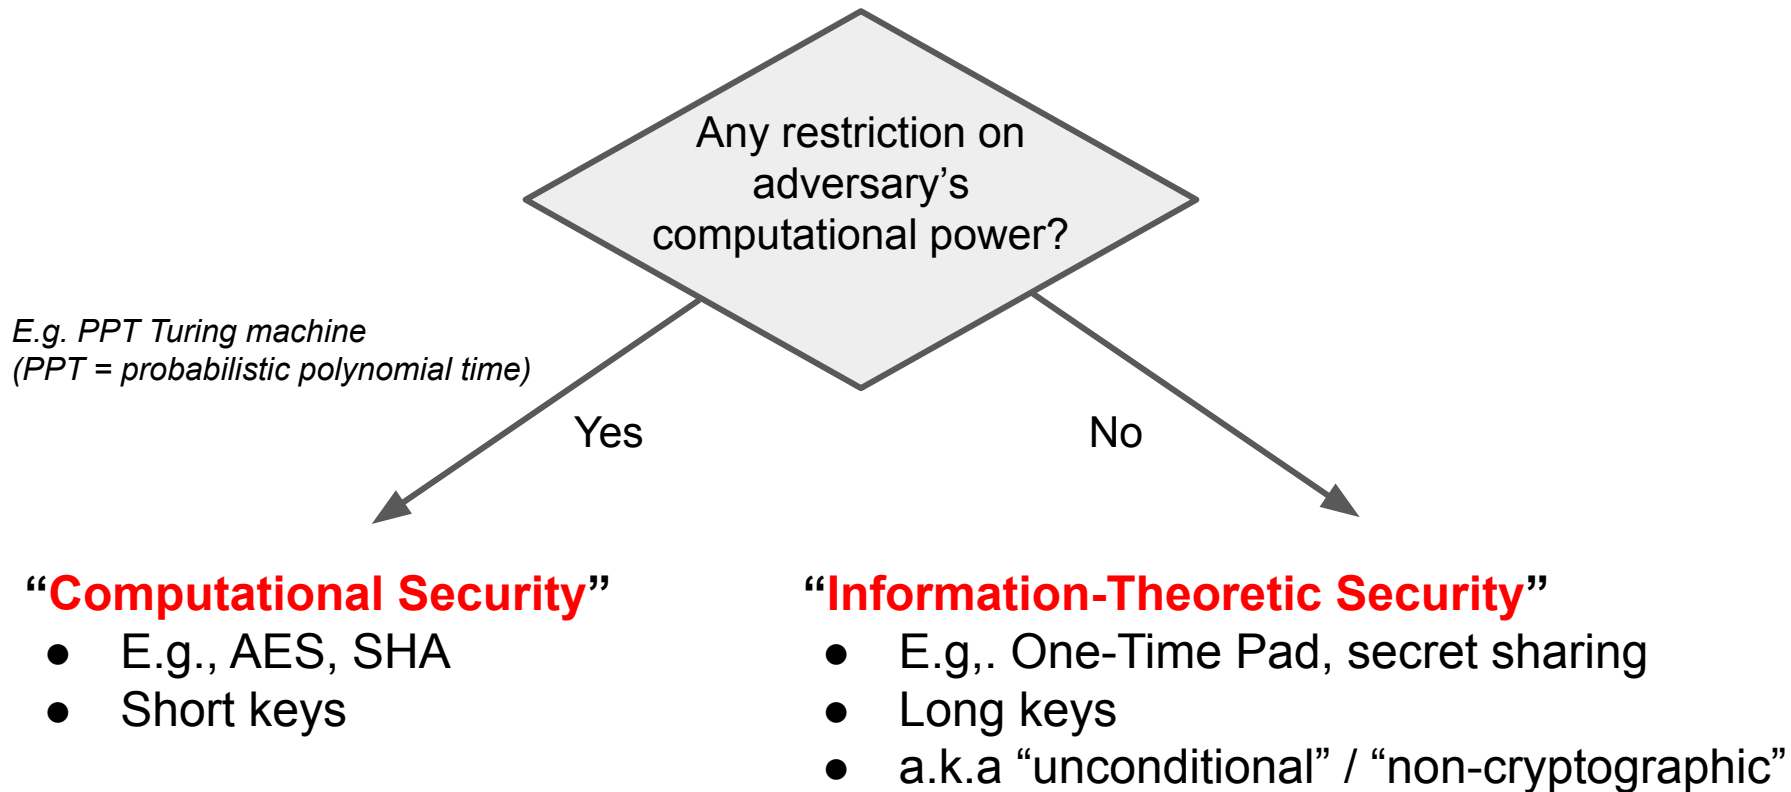
└ Where do I get one of those?

The One-Time Pad

$$\begin{array}{rcl} & 1 & 1 & 0 & 0 & 1 & 0 & \text{message} \\ \oplus & 0 & 1 & 0 & 1 & 0 & 1 & \text{key} \\ \hline & 1 & 0 & 0 & 1 & 1 & 1 & \text{ciphertext} \end{array} \quad \left. \vphantom{\begin{array}{rcl} & 1 & 1 & 0 & 0 & 1 & 0 & \text{message} \\ \oplus & 0 & 1 & 0 & 1 & 0 & 1 & \text{key} \\ \hline & 1 & 0 & 0 & 1 & 1 & 1 & \text{ciphertext} \end{array}} \right\} \text{store these}$$

- 😊 Reveals **no information** about the message (“perfect secrecy”)
- 😞 Key **must be** as long as the message
- 😞 Loss of key or ciphertext \Rightarrow loss of message

Two Security Notions



Does Computational Security Exist?

Impagliazzo's Five Worlds

Algorithmica
($P=NP$)

Heuristica
(NP easy on average)

Pessiland
(NP hard on average + no OWFs)

No

Minicrypt
(OWFs exist, but no public key crypto)

Cryptomania
(public key crypto exists)

Yes

Which world do we live in???

OWF = One-Way Function

Cryptographic Obsolescence

DES, SHA-1, MD5



*Using Shor's
algorithm

{ Discrete Logarithm Problem
Factoring



“Hope that whoever you’re trying to keep the secret from is not a better mathematician than you are” - Michael Sipser

Computationally Secure Archives

“Cascade cipher”

$$\text{AES-256}_{k1}(\text{Blowfish}_{k2}(\text{3DES}_{k3}(m)))$$

Used in ArchiveSafeLT (Sabry & Samavi. ACSAC'22)

“All-or-nothing Transform”

Parse m as $m_1m_2m_3$

$$c_i := m_i \oplus \text{AES-256}_k(i+1) \text{ for } i=1,2,3$$

$$c_4 := k \oplus \text{SHA-256}(c_1, c_2, c_3)$$

$$\text{Reed-Solomon}(c_1, c_2, c_3, c_4)$$

Used in AONT-RS (Resch & Plank. FAST'11)

Regardless – computational security is susceptible to Harvest Now, Decrypt Later

Additive Secret Sharing

What if we generalized the One-Time Pad?

$$m \oplus r_1 \oplus r_2 \oplus \dots \oplus r_{n-1} =: c$$

Give each of these “shares” to a
different party

Seems even more useless than One-Time Pad:

- Need n parties to store each “share” (and each share has same size as m)
- Cannot tolerate loss of a single share // what if we fixed this issue?

Shamir's Secret Sharing

Goal: share a secret m among n parties such that any $t < n$ parties can reconstruct m , but fewer than t parties learn **no information** about m .


Given m (as a finite field element), and integers $t < n$, do the following:

1. Pick $t-1$ random field elements r_1, \dots, r_{t-1}
2. Define polynomial $p(x) = m + r_1x + \dots + r_{t-1}x^{t-1}$
3. Pick n arbitrary non-zero field elements x_1, \dots, x_n
4. Return shares $(x_1, p(x_1)), \dots, (x_n, p(x_n))$

Given any t shares, can uniquely interpolate p and retrieve $m = p(0)$

Just a non-systematic Reed-Solomon code in disguise...


Secret Shared Archives



- 😊 Shamir's secret sharing provides information-theoretic "perfect secrecy" (like OTP)
- 😊 Tolerates lost/stolen shares
- 😞 High storage overhead
- 😞 Given time, adversary may eventually steal a threshold number of shares


Used in many secure archival works:
POTSHARDS, PASIS, LINCOS, etc.

Proactive Secret Sharing (PSS)




- 😊 Each party gets a new share independent of its old share
- 😊 Old/stolen shares obsolete *only if* honest parties delete their shares
- 😊 Nobody learns anything other than their new share
- 😊 Byzantine fault tolerant
- 😞 High communication overheads

Data in Transit



- May be easier for adversary to eavesdrop data in transit
- TLS encryption is only computationally secure!
- Can we protect data in transit information-theoretically?

Data in Transit



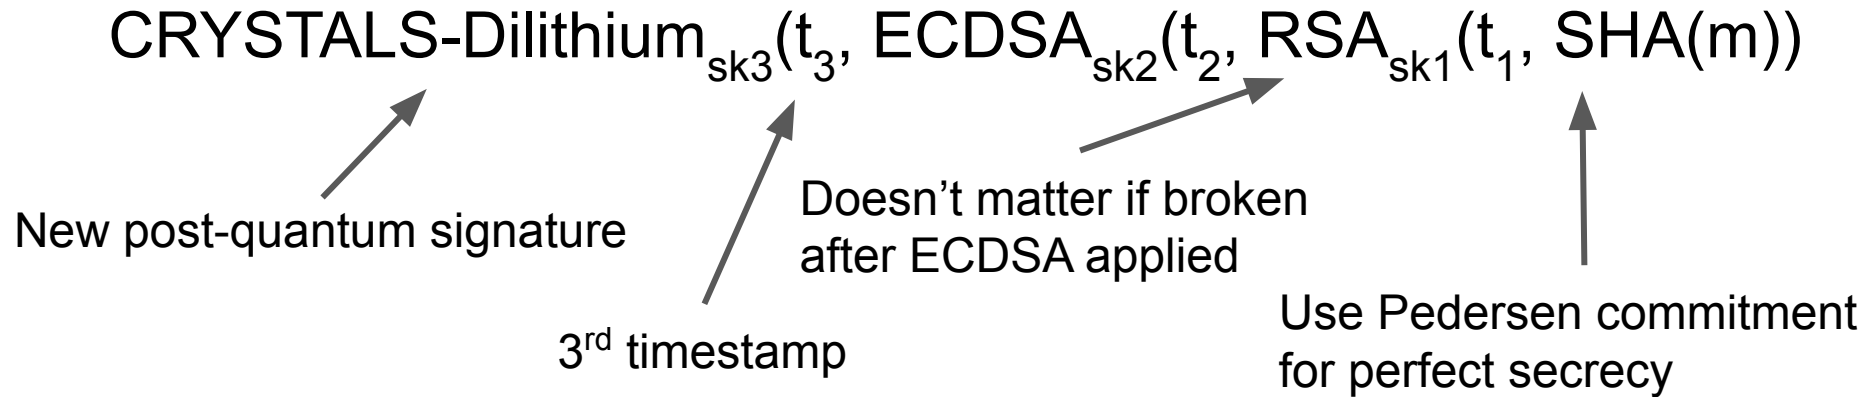
Yes: Quantum Key Distribution (QKD)

- 😊 Can generate a OTP key in a way impervious to eavesdropping and tampering
- 😞 Specialized infrastructure
- 😞 Not practical yet

Used in LINCOS (Braun et al. ASIA CCS'17)

Integrity

- Normally we use digital signatures for integrity against malicious tampering
- But digital signatures are also susceptible to cryptographic obsolescence
- **Solution**: use a chain of digitally signed timestamps

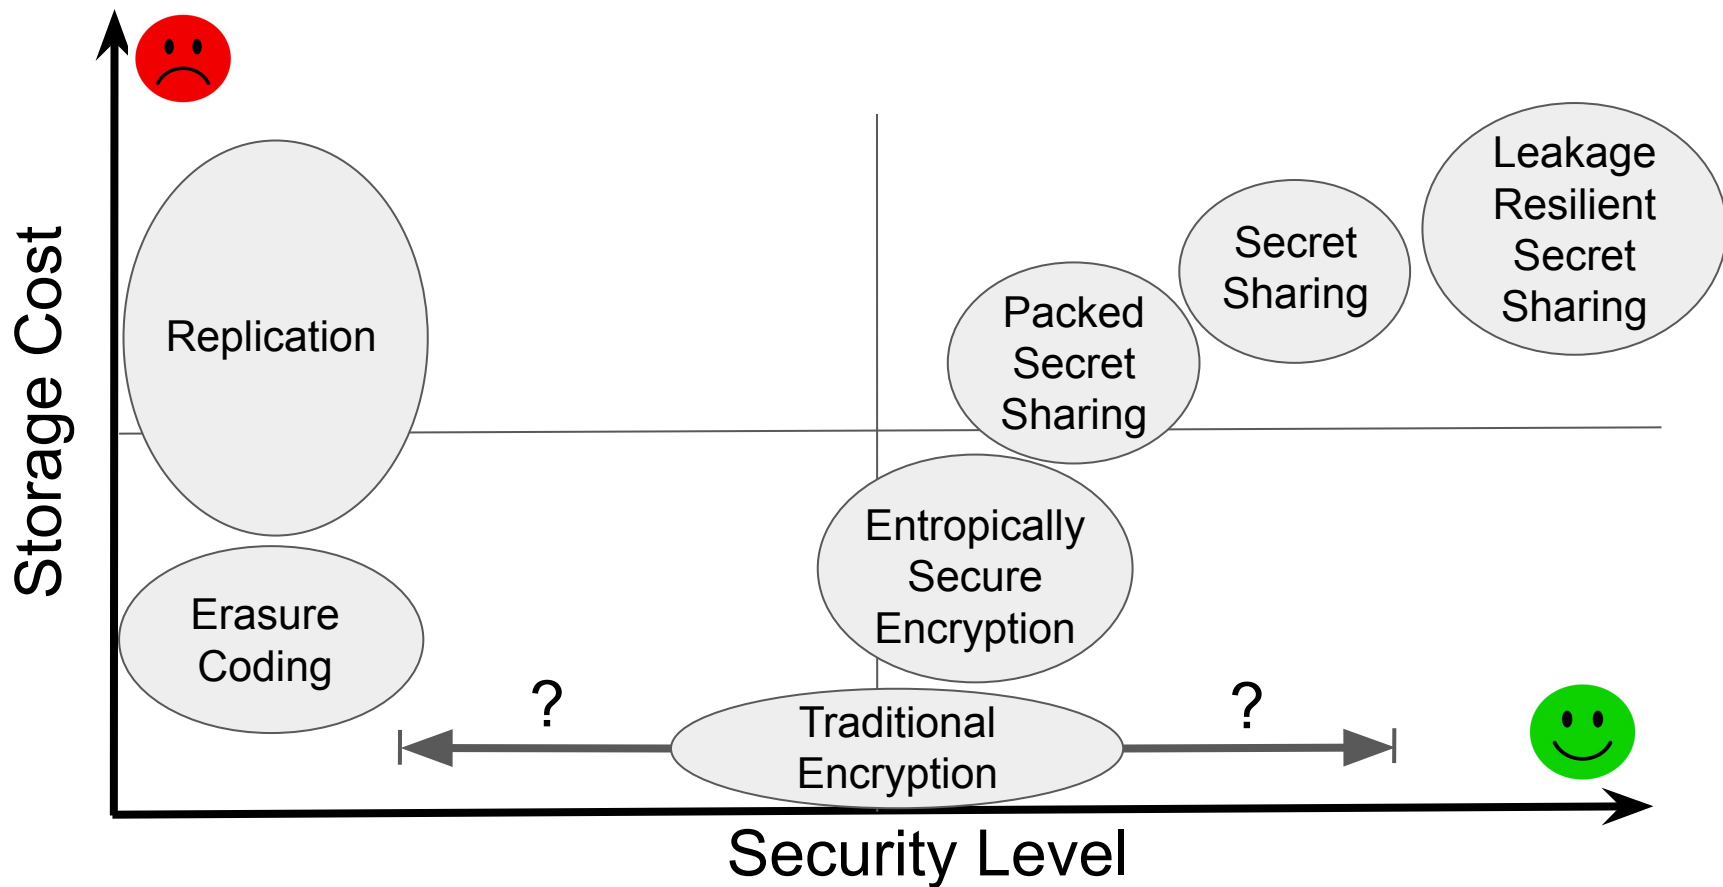


We're all caught up

Systems	Confidentiality		Storage Cost
	In Transit	At Rest	
ArchiveSafeLT	Computational	Computational	Low
AONT-RS	Computational	Computational	Low
HasDPSS	Computational	ITS	High
LINCOS	ITS	ITS	High
PASIS	Computational	ITS (sometimes)	Low-High
POTSHARDS	Computational	ITS	High
VSR Archive	Computational	ITS	High
AWS, Azure, Google Cloud	Computational	Computational	Low

What can be improved?

A Fundamental Trade-off



Make Storage Cheaper

Glass



😊 Dense ($\sim 429 \text{ TB/in}^3$)

😊 Survives millenia

😊 No maintenance

DNA



😊 Super dense ($\sim 1 \text{ EB/mm}^3$)

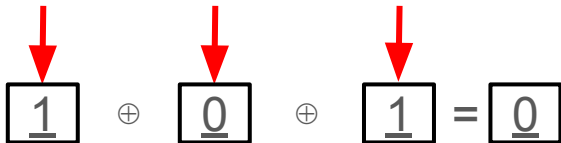
😊 Survives centuries

😞 DNA synthesis slow + costly

Leakage-Resilient Secret Sharing

Simple leakage attack on 3-of-3 additive secret sharing

$$\boxed{101} \oplus \boxed{010} \oplus \boxed{011} = \boxed{100}$$


$$\boxed{1} \oplus \boxed{0} \oplus \boxed{1} = \boxed{0}$$


Leakage attack: adversary may leak a few bits of information about each secret share undetected via hidden side-channel.

Leakage-resilient secret sharing (LRSS) to the rescue?

Questions:

- Proactive LRSS?
- What is the right leakage model?

Bounded Storage Model (BSM)



BSM in a nutshell:

- Assume restrictions on adversarial **storage capacity**, and use a lot of public randomness.
- Yields information-theoretic channels.
- overdue for practical evaluation



Q&A

Secure Archival is Hard... *Really* Hard

Christopher Smith, Maliha Tabassum, Soumya Chowdary Daruru,
Gaurav Kulhare, Arvin Wang, Ethan L. Miller^{*†}, Erez Zadok

Stony Brook University, PureStorage^{}, UC Santa Cruz[†]*

