Understanding the Threat Landscape of Remote Cloud Storage

ECE239AS

Slides created by Katherine Izhikevich

“Origins” of a Data Breach

Zero Day

Phishing Attack

Stolen Credential

“Origins” of a Data Breach

Zero Day

Phishing Attack

Stolen Credential

Unintentional Exposure

- Using a resource that seems to work as intended, but is vulnerable in practice.
- Think: using `1234` as a password.

Data leaks happen all the time

Aston Villa Football Club leaks members' data

21 May 2024 Chris Price Cybersecurity

Aston Villa Football Club (AVFC) left a publicly leaking Amazon Web Services (AWS) S3 bucket containing the personally identifiable information of 135,770 individuals. The affected fans are vulnerable to spear phishing, social engineering attacks, and identity theft attempts.

Data leaks happen all the time

India's national logistics portal exposed sensitive personal data, trade records

Jagmeet Singh / 4:45 AM PDT • October 2, 2023

"... Exposed data included full names, nationality, date of birth, gender, passport numbers, ..."

Data leaks happen all the time

Lantum S3 bucket leak is prescription for chaos for thousands of UK doctors

Lindsay Clark Mon 12 Jun 2023 // 12:34 UTC

"... exposed personal details relating to 3,200 individuals via unsecured S3 buckets… including passport details… medical documents …

Data leaks happen all the time

A Cyberattack Illuminates the Shaky State of Student Privacy

At a moment when education technology firms are stockpiling sensitive information on millions of school children, safeguards for student data have broken down.

By Natasha Singer

July 31, 2022

6 MIN READ

Data leaks happen all the time

A Cyberattack Illuminates the Shaky State of Student Privacy

At a moment when education technology firms are stockpiling sensitive information on millions of school children, safeguards for student data have broken down.

By Natasha Singer
July 31, 2022

6 MIN READ

Microsoft leaked 2.4TB of data belonging to sensitive customer. Critics are furious

Data includes signed contracts and projects related to critical infrastructure.

DAN GOODIN - 10/20/2022, 4:03 PM

Data leaks happen all the time

A Cyberattack Illuminates the Shaky State of Student Privacy

At a moment when education technology firms are stockpiling sensitive information on millions of school children, safeguards for student data have broken down.

By Natasha Singer
July 31, 2022

6 MIN READ

Microsoft leaked 2.4TB of data belonging to sensitive customer. Critics are furious

Data includes signed contracts and projects related to critical infrastructure.

DAN GOODIN - 10/20/2022, 4:03 PM

A huge data leak of 1 billion records exposes China's vast surveillance state

One billion resident records were allegedly siphoned from a police database

Zack Whittaker, Carly Page / 1:15 PM EDT • July 7, 2022

Comment

Data leaks happen all the time

A Cyberattack Illuminates the Shaky State of Student Privacy

At a moment when education technology firms are stock[illegible] millions of school children, safeguards for student data have broken down.

By Natasha Singer

July 31, 2022

6 MIN READ

Micros[illegible] data belonging to ser[illegible] critics are furious

[illegible] related to critical infrastructure.

Insecure Amazon S3 bucket exposed personal data on 500,000 Ghanaian graduates

John Leyden 06 January 2022 at 10:58 UTC

Updated: 10 January 2022 at 09:40 UTC

A[illegible] leak of 1 billion records exposes C[illegible] vast surveillance state

One[illegible]ion resident records were allegedly siphoned from a police database

Zack Whittaker, Carly Page / 1:15 PM EDT • July 7, 2022

Comment

Data leaks happen all the time

Data leaks happen all the time

Data leaks happen all the time

What is a bucket?

File storage system in the cloud

What is a bucket?

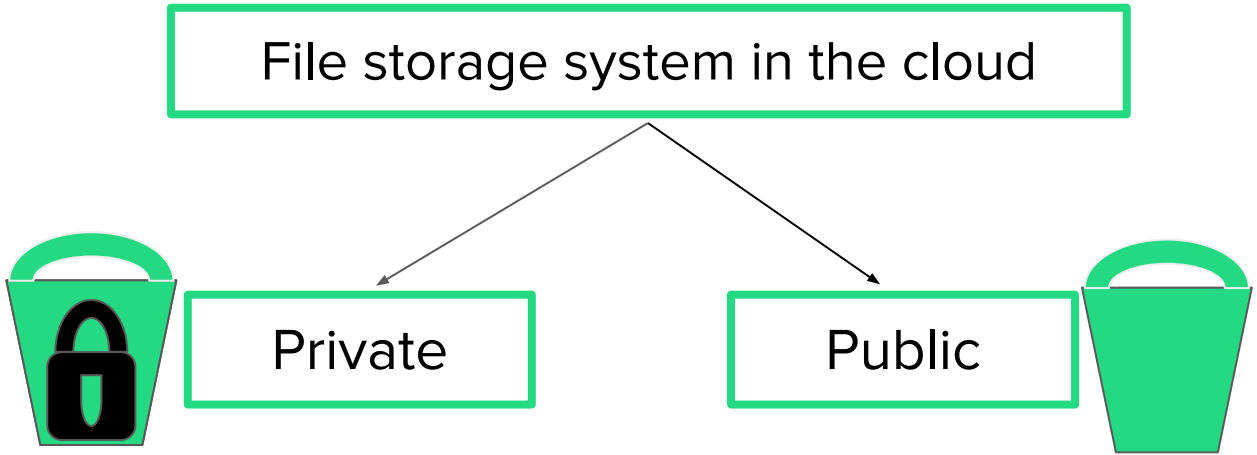

What is a bucket?

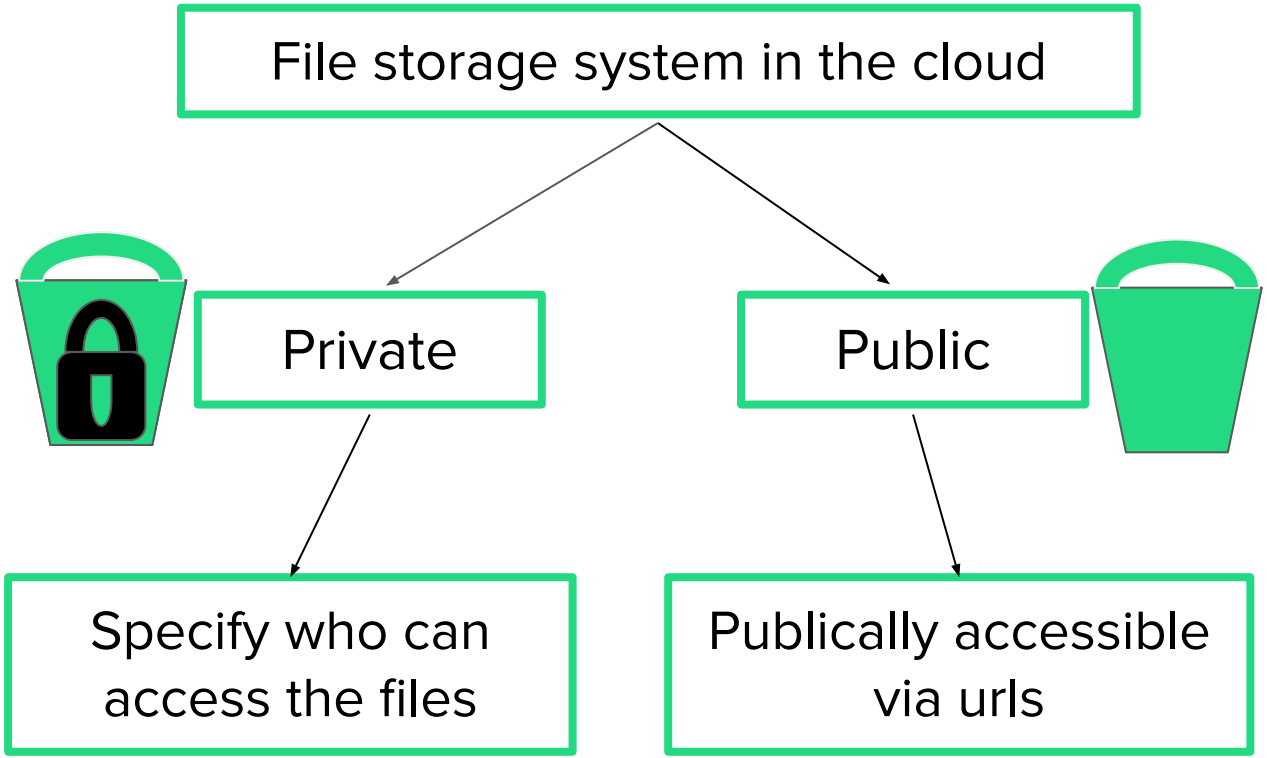

Private buckets store sensitive information

File storage system in the cloud

Private

Specify who can access the files

Storing sensitive information about customers

Private buckets store sensitive information

File storage system in the cloud

Private

Specify who can access the files

Employees, contractors

Storing sensitive information about customers

Public buckets host public information

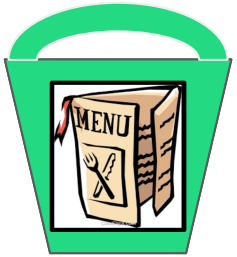

File storage system in the cloud

Public

Publically accessible via urls

Public buckets host public information

File storage system in the cloud

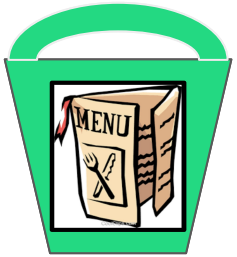

Public

https://**myrestaurant**.s3.us-west-2.amazonaws.com

Publically accessible via urls

Public buckets can be configured as a website

https://**mywebsite**.s3.us-west-2.amazonaws.com

Publically accessible via this url

Creative Freedom to Name Buckets

- Alphanumeric, lowercase only, dashes allowed
- Length 3-63
- Must be an unused/unique name - no repository exists

Creative Freedom to Name Buckets

- Alphanumeric, lowercase only, dashes allowed
- Length 3-63
- Must be an unused/unique name - no repository exists
- Examples of valid bucket names (that are available)!
  - thebigtoe
  - 3140ebu3b
  - you-get-the-point

Creative Freedom to Name Buckets

- Alphanumeric, lowercase only, dashes allowed
- Length 3-63
- Must be an unused/unique name - no repository exists
- Examples of valid bucket names (that are available)!
  - thebigtoe
  - 3140ebu3b
  - you-get-the-point
- Astronomical number of possible bucket names: ~10^{101}

Cloud Providers Offering Bucket Service

Cloud Providers Offering Bucket Service

Simple Storage Service

Bucket Recap

File storage system in the cloud with a unique name

Private

Public

Specify who can access the files

Publically accessible via urls

Attackers exploit public buckets meant to be private

File storage system in the cloud with a unique name

Private

Public

Specify who can access the files

Publically accessible via urls

Researchers explore “leaky” cloud infrastructure

ACSAC ‘18

There’s a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Seminal study, guessing bucket names, finding exposed data in s3

RAID ‘21

Stratosphere: Finding Vulnerable Cloud Storage Buckets

Jack Cable* — Stanford University
Drew Gregory* — Stanford University
Liz Izhikevich* — Stanford University
Zakir Durumeric — Stanford University

Faster, ML-based bucket name generator across s3, GCP, and Alibaba

EuroS&P ‘24

Using Honeybuckets to Characterize Cloud Storage Scanning in the Wild

Katherine Izhikevich — *UC San Diego*
Geoffrey M. Voelker — *UC San Diego*
Stefan Savage — *UC San Diego*
Liz Izhikevich — *Stanford University*

Deploying honeypots to understand threat landscape of exposures

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella*
UC Santa Barbara
conand@cs.ucsb.edu

Mario Polino
Politecnico di Milano
mario.polino@polimi.it

Marcello Pogliani
Politecnico di Milano
marcello.pogliani@polimi.it

Stefano Zanero
Politecnico di Milano
stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella*
UC Santa Barbara
conand@cs.ucsb.edu

Mario Polino
Politecnico di Milano
mario.polino@polimi.it

Marcello Pogliani
Politecnico di Milano
marcello.pogliani@polimi.it

Stefano Zanero
Politecnico di Milano
stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

$

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella*
UC Santa Barbara
conand@cs.ucsb.edu

Mario Polino
Politecnico di Milano
mario.polino@polimi.it

Marcello Pogliani
Politecnico di Milano
marcello.pogliani@polimi.it

Stefano Zanero
Politecnico di Milano
stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

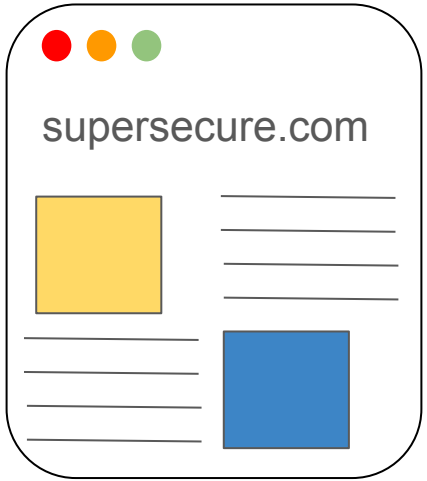

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

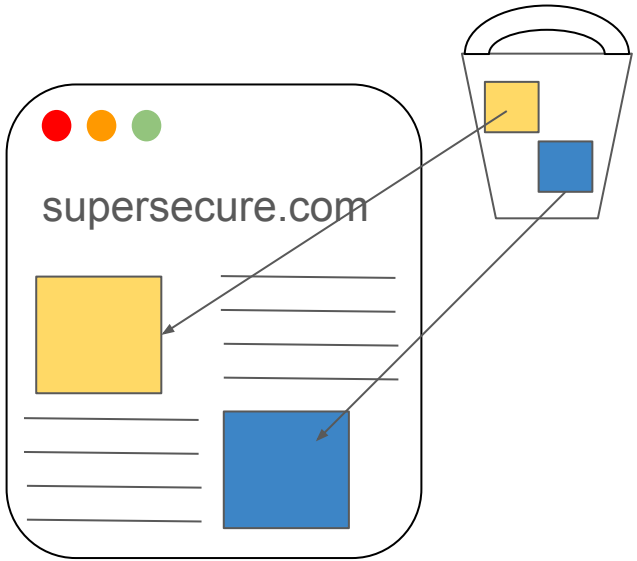

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

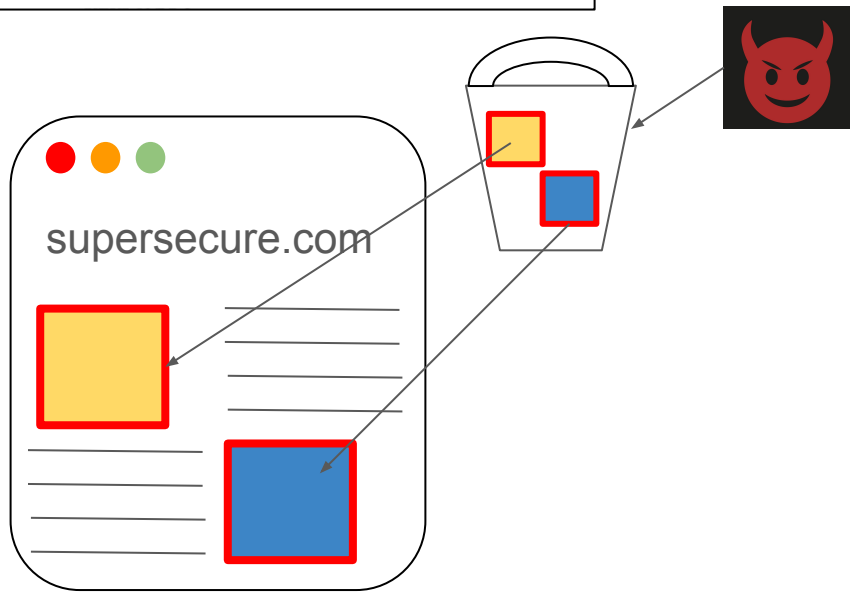

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

supersecure.com

assets.supersecure.com

↓

assets-supersecure.s3.amazonaws.com

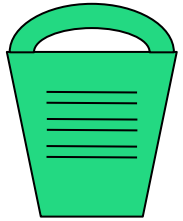

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

supersecure.com

assets.supersecure.com

↓

assets-supersecure.s3.amazonaws.com

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

supersecure.com

assets.supersecure.com

assets-supersecure.
s3.amazonaws.com

The Bucket Threat Model

There's a Hole in that Bucket!
A Large-scale Analysis of Misconfigured S3 Buckets

Andrea Continella* — UC Santa Barbara — conand@cs.ucsb.edu
Mario Polino — Politecnico di Milano — mario.polino@polimi.it
Marcello Pogliani — Politecnico di Milano — marcello.pogliani@polimi.it
Stefano Zanero — Politecnico di Milano — stefano.zanero@polimi.it

Threat Model

Readable Buckets	- *Data leakage*
Writable Buckets	- *Ransom Demand* - *Web Resource Infection*
Unclaimed Buckets	- *Dangling Subdomain Takeover*

supersecure.com

assets.supersecure.com

assets-supersecure.
s3.amazonaws.com