Cloud Security

S.Venkatesan

Acknowledgement: The contents, example scripts and some figures are copied from various sources.
Thanks to all authors and sources made those contents public and usable for educational purpose

NIST reference architecture

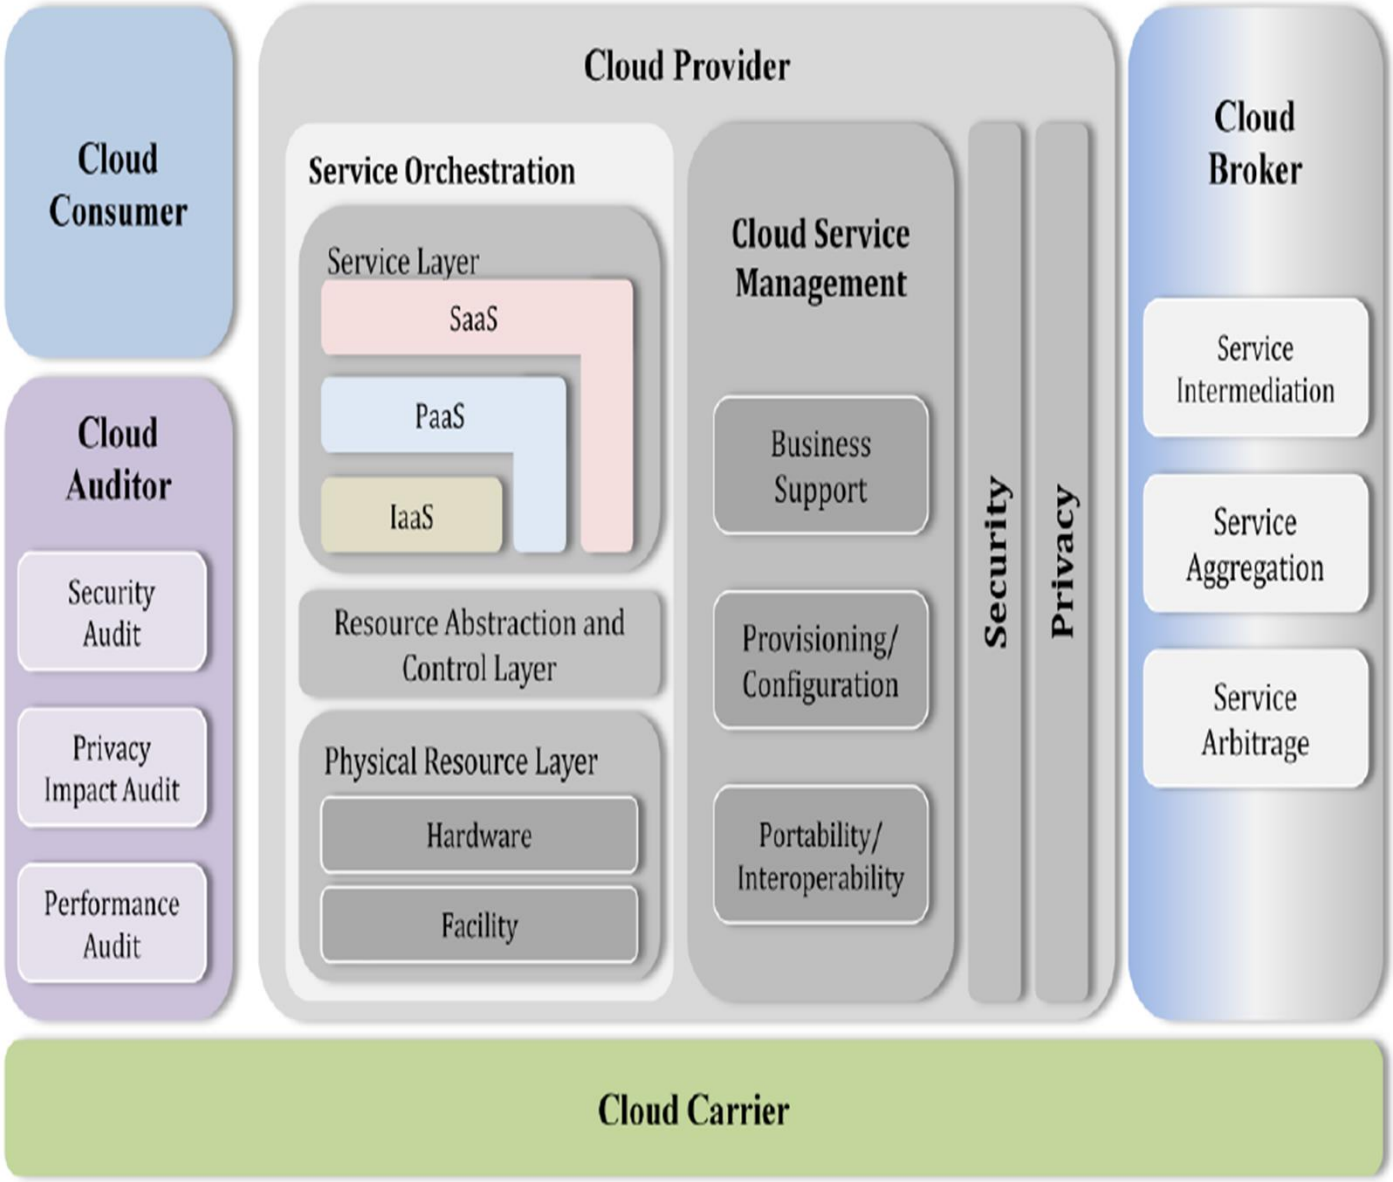

Security Requirements

Cloud-based control approach

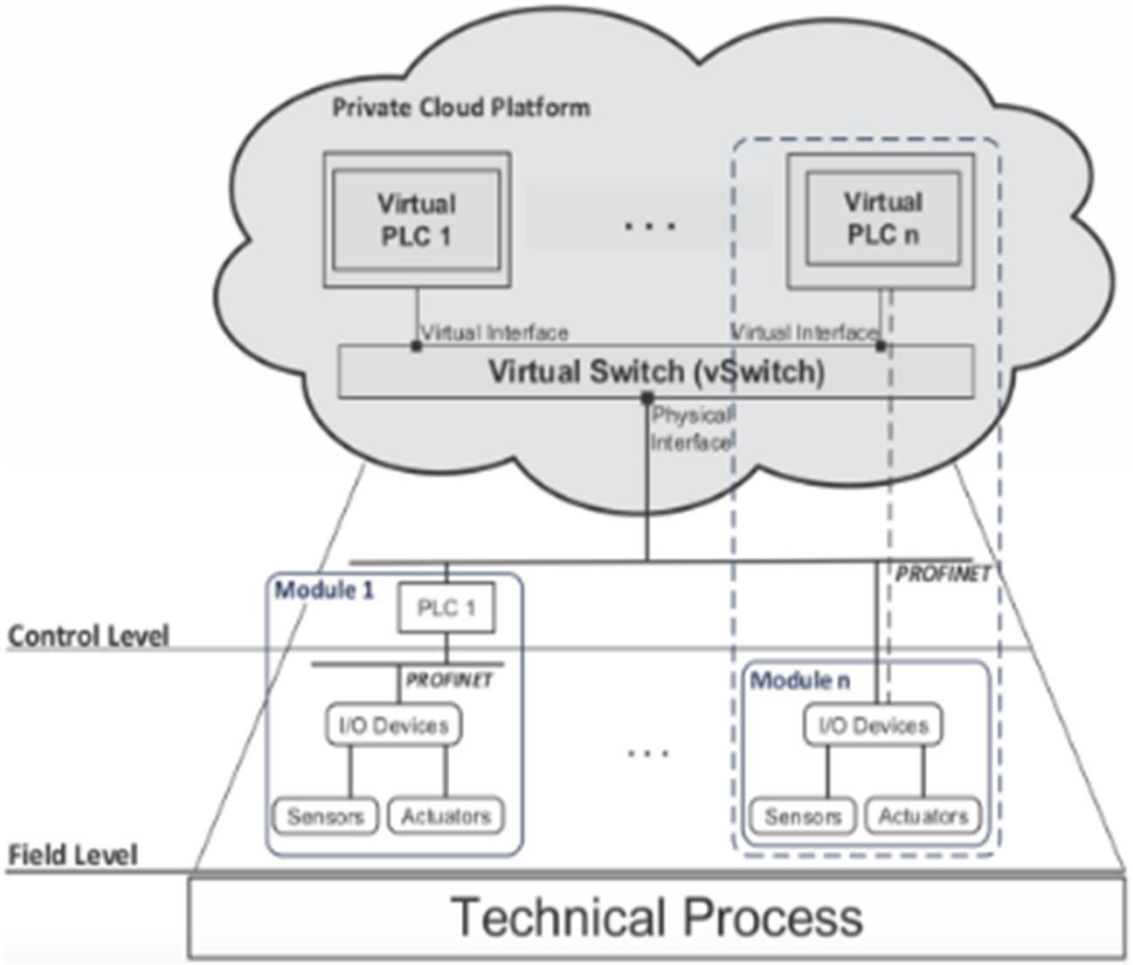

Statistics

- 92% Enterprises Have a Multi-Cloud Strategy in Place
- 98% of Organizations Experienced a Cloud Security Breach in Past 18 Months
- 72% of IT Security Leaders Rank Cloud as Top Digital Transformation Priority
- ~90% of Data Breaches Target Servers
- 96% of Web App Attack-Based Mail Server Compromises Involve the Cloud

TPM

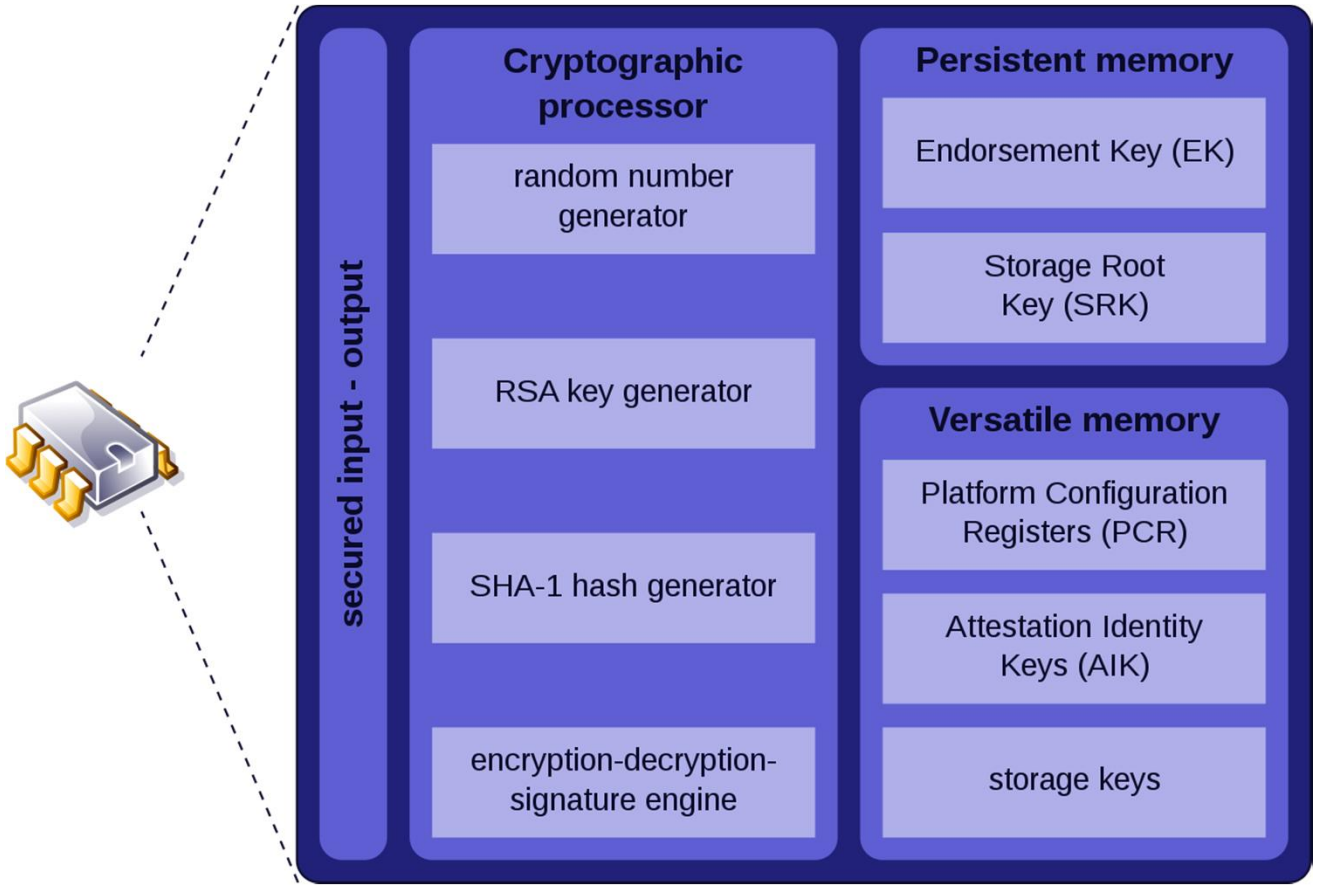
secured input - output
Cryptographic processor
random number generator
RSA key generator
SHA-1 hash generator
encryption-decryption-signature engine
Persistent memory
Endorsement Key (EK)
Storage Root Key (SRK)
Versatile memory
Platform Configuration Registers (PCR)
Attestation Identity Keys (AIK)
storage keys

Google Cloud Security

- Communications over the internet to Google public cloud services are encrypted in transit.
- To protect DoS Google Cloud Armor
- Identities, users, and services are strongly authenticated. Access to sensitive data is protected by advanced tools like phishing-resistant security keys (**Titan Security Key)**.
- Data stored on Google infrastructure is automatically encrypted at rest and distributed for availability and reliability.

Layers of encryption

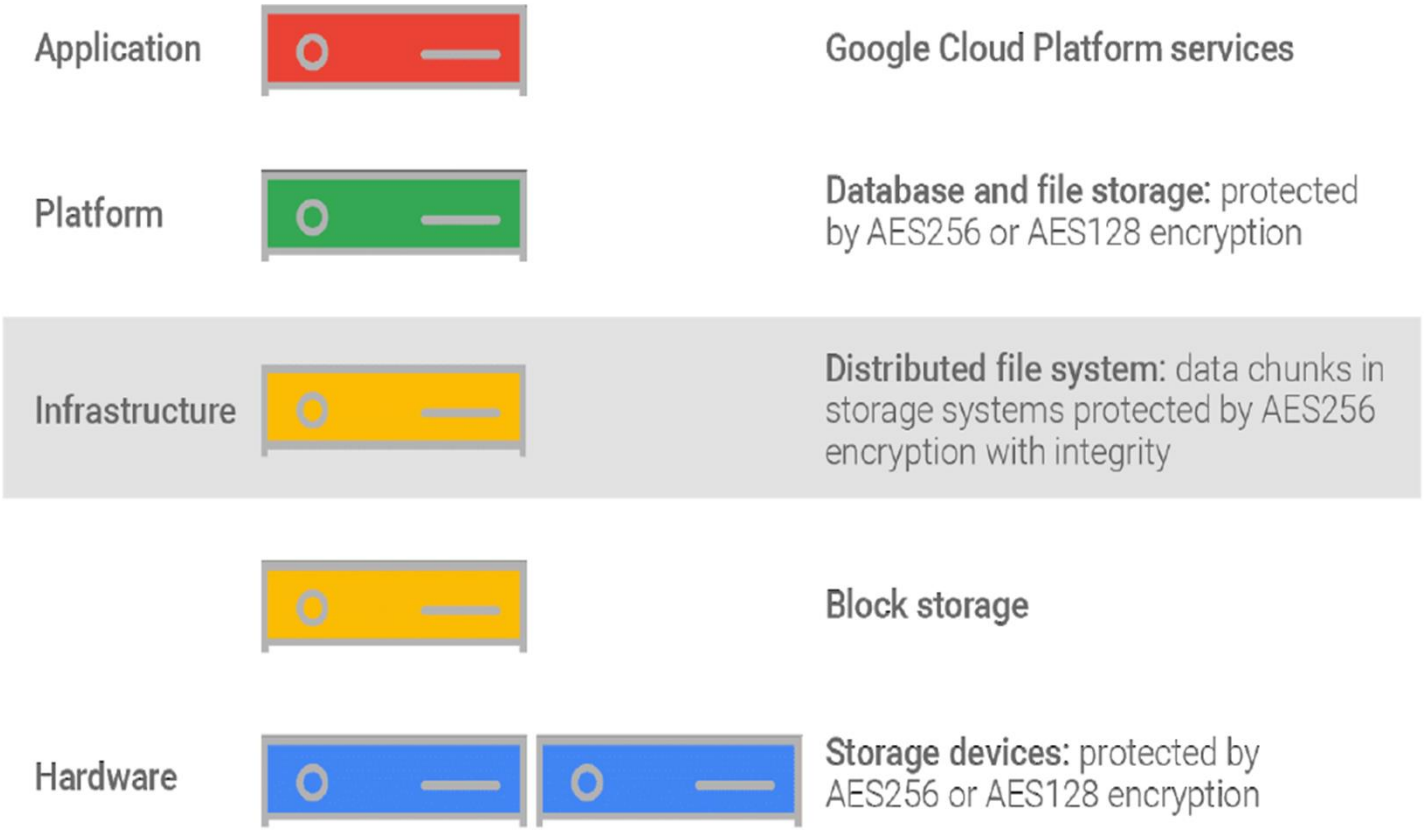

Data at Google is broken up into encrypted chunks for storage& Key Hierarchy

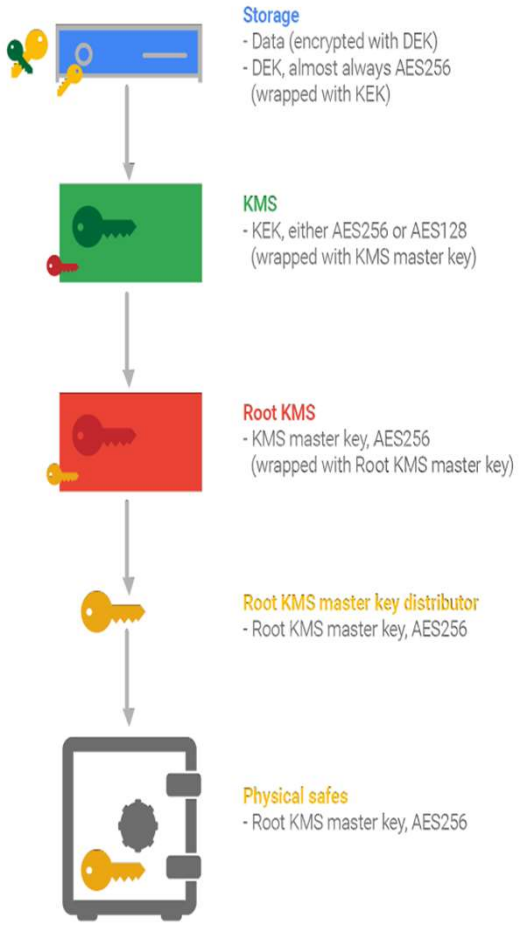

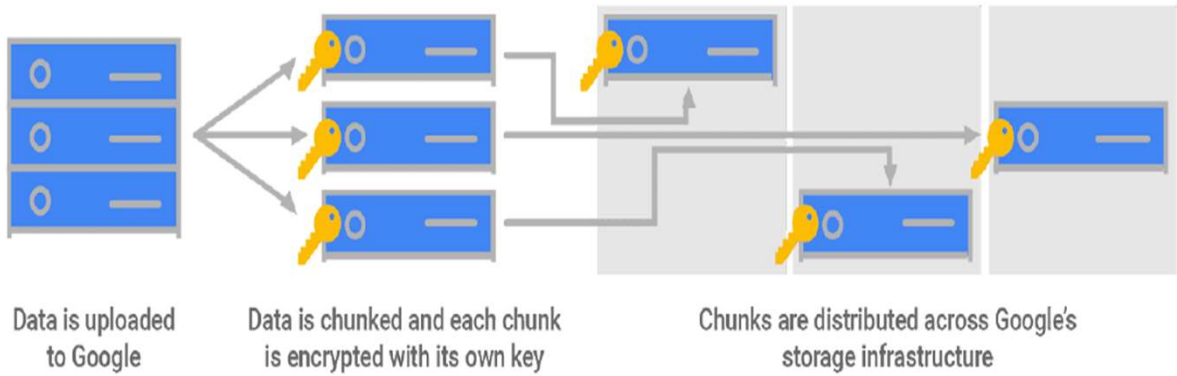

Protecting Data at Rest on Amazon S3 Glacier

- Data at rest stored in Amazon S3 Glacier is automatically server-side encrypted using 256-bit Advanced Encryption Standard (AES-256) with keys maintained by AWS.

- The encryption key is then encrypted itself using AES-256 with a master key that is stored in a secure location.

- The master key is rotated on a regular basis.

Protecting Data at Rest on Amazon RDS

- INSERT INTO Customers (CustomerFirstName,CustomerLastName) VALUES (AES_ENCRYPT('John',@key), AES_ENCRYPT('Smith',@key);

Harden the VMs – Google Cloud

- Integrity of the VMs will be checked.
- Secure Boot
- Virtual Trust Platform Module (TPM) - Platform integrity, Disk encryption, Password protection.
- Google Titan Chip (similar to TPM) **-** securely identify and authenticate legitimate access at the hardware level, minimizing the chances of running altered software.

Harden Your VMs with Shielded Computing

Security Assurances

- Malicious guest OS firmware, including malicious UEFI drivers
- Malicious guest OS, including boot and kernel vulnerabilities
- Malicious insiders within your organization

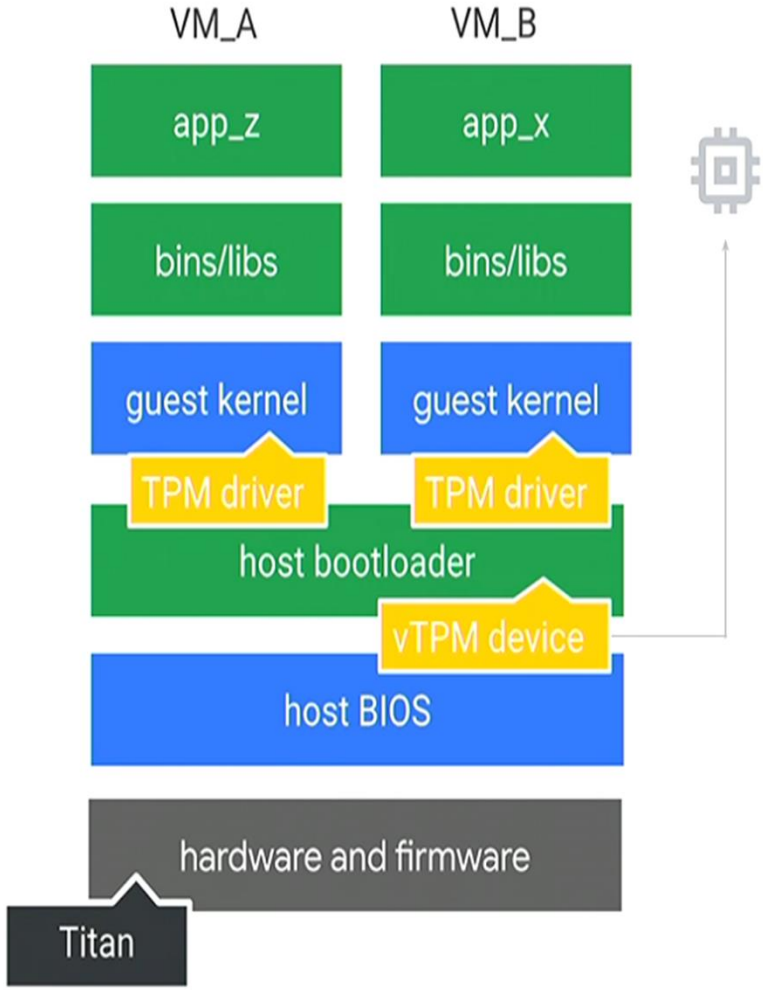

Google Cloud

Integrity of the Data

- One of the problems associated with outsourcing data to cloud service providers is the data integrity of outsourced data.

- Data integrity encompasses the
  - Completeness
  - Correctness and
  - Freshness

Merkle Tree

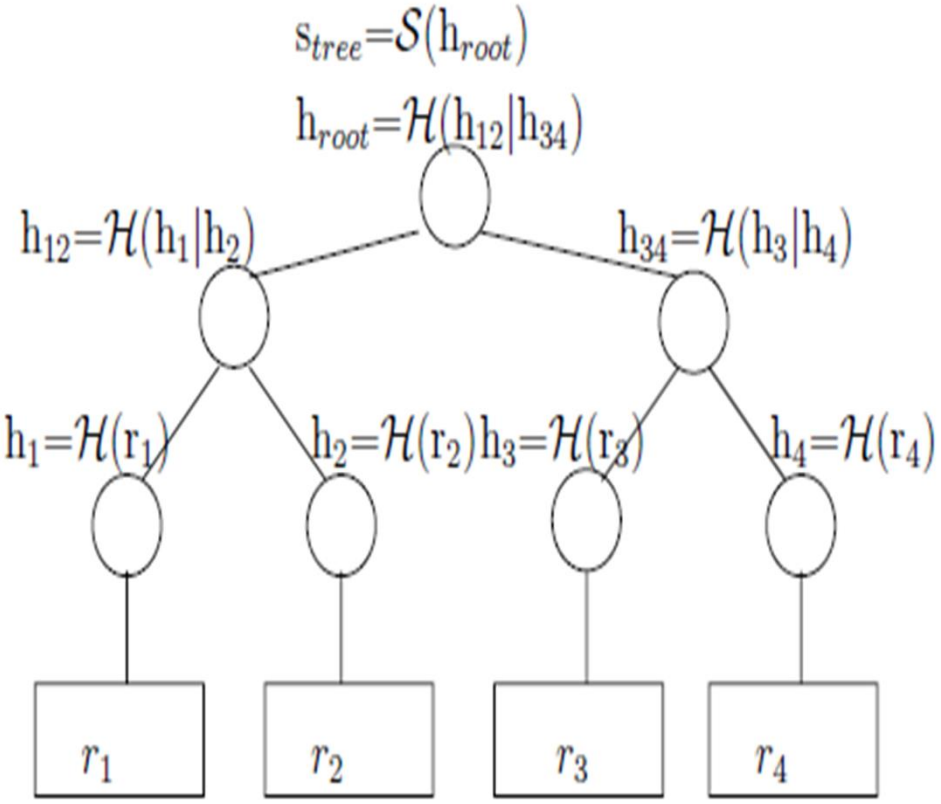

Feifei Li, Marios Hadjieftheriou, George Kollios, and Leonid Reyzin, Authenticated Index Structures for Outsourced Databases Handbooks on Database Security

An Example

Table: Employee table

ID	Name	Salary
10	Alice	1,000
20	Bob	2,000
30	Cindy	3,000
40	Dan	4,000
50	Eva	5,000
60	Felix	6,000

Merkle Tree

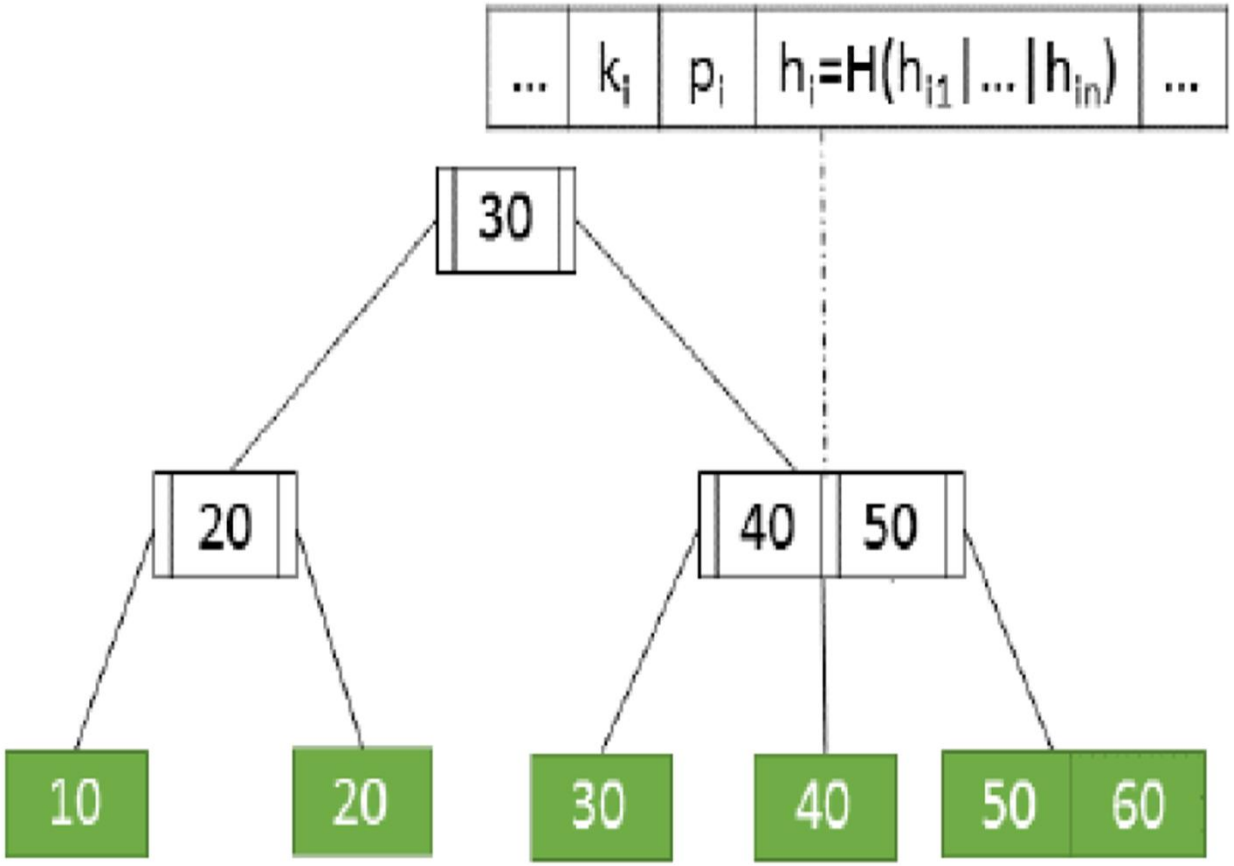

Radix Path Identifier

- l - level of the node
- r_b - radix base
- f - fanout of the MHT
- i - index of the pointer

Calculation of Radix Path Identifier

$$rpi = \begin{cases} l & \text{if } l == 0 \\ rpi_{parent} * r_b + i & \text{if } l > 0 \end{cases}$$

Merkle Hash Tree with Radix Path Identifiers

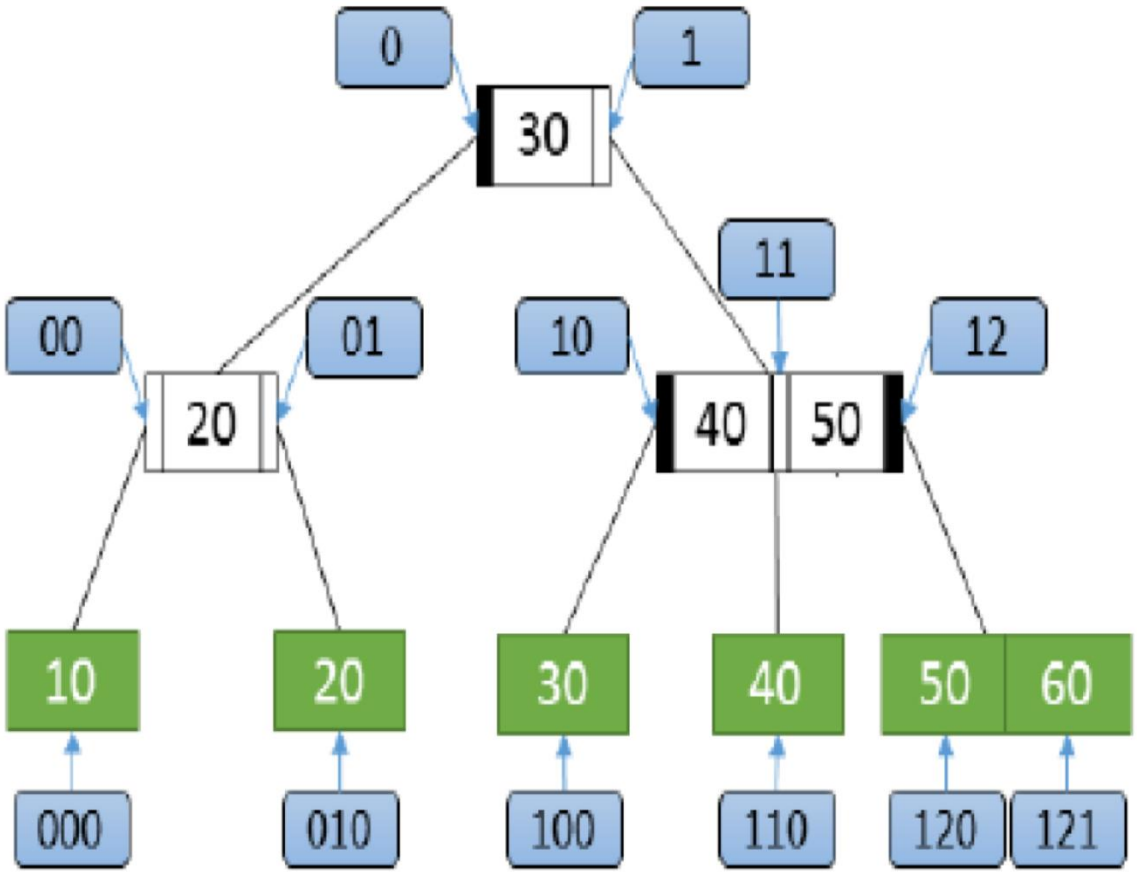

Properties

- RPIs are continuous in nodes, but not continuous among two consecutive nodes.
- From an RPI, we can easily find the RPI of its parent pointer based on the fact that rpi_{parent} equals to $\lfloor rpi/rb \rfloor$.
- From the RPI in a node, we can easily calculate the min and max RPIs in the node, which are $(\lfloor rpi/rb \rfloor) * r_b + (r_b - 1)$.
- From an RPI in a node, we can easily comute the index i of the pointer or key in the node, which is rpi mod r_b.

Single Authentication Table

ID	RPI	Hash	Level
-1	0	hash	2
30	1	hash	2
-1	0	hash	1
20	1	hash	1
-1	3	hash	1
40	4	hash	1
50	5	hash	1
10	0	hash	0
20	3	hash	0
30	9	hash	0
40	12	hash	0
50	15	hash	0
60	16	hash	0

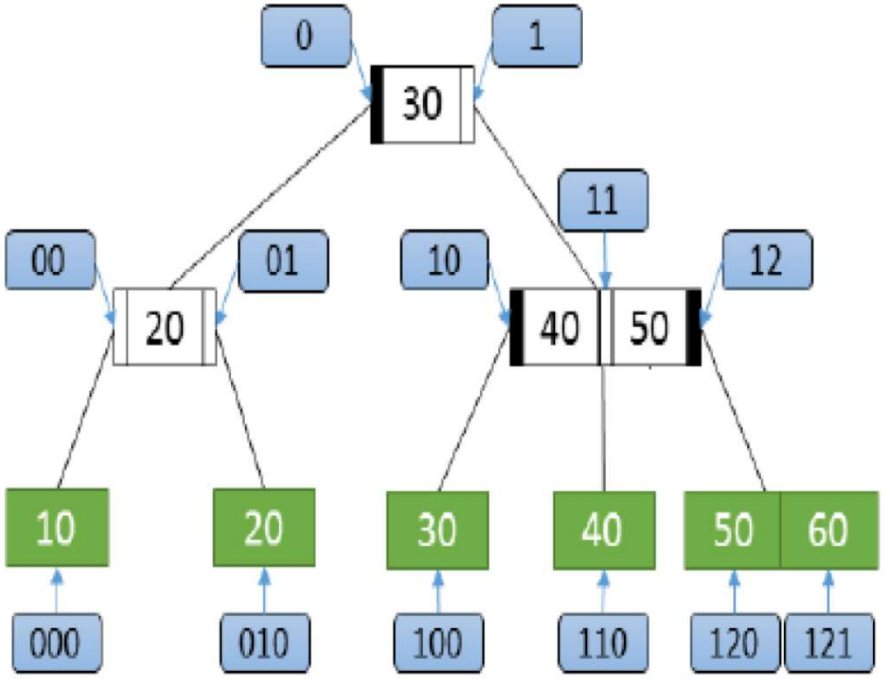

Level Based Authentication Table

Emp_2 (Root)		
ID	**RPI**	**Hash**
-1	0	hash
30	1	hash

Emp_1		
ID	**RPI**	**Hash**
-1	0	hash
20	1	hash
-1	3	hash
40	4	hash
50	5	hash

Employee (Leaf Node)				
ID	**Name**	**Salary**	**RPI**	**Hash**
10	Alice	1000	0	hash
20	Bob	2000	3	hash
30	Cindy	3000	9	hash
40	Dan	4000	12	hash
50	Eva	5000	15	hash
60	Felix	6000	16	hash

Multi-Join Query

```
select a0.RPI as RPI0, a0.hash as hash0, a1.RPI as RPI1, a1.hash as
hash1, a2.RPI as RPI2, a2.hash as hash2 from Employee emp
left join Employee a0 on a0.RPI/3 = emp.RPI/3
left join Emp_2 a1 on a1.RPI/3 = emp.RPI/(3*3)
left join Emp_2 a2 on a2.RPI/3 = emp.RPI/(3*3*3)
where emp.ID = 40;
```

References

- Erik Kajati, Peter Papcun, Chao Liu, Ray Y. Zhong, Jiri Koziorek, Iveta Zolotova, Cloud based cyber-physical systems: Network evaluation study, Advanced Engineering Informatics, Volume 42, 2019.

- https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication500-292.pdf

- https://www.apriorit.com/dev-blog/545-sandbox-evading-malware