International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 10 | Oct 2026 www.irjet.net p-ISSN: 2395-0072

A Systematic Comparative Review of Cloud Ransomware Recovery Techniques: Data Integrity, Recovery Performance, and Cost

Dr. Meghendra Kumar Gajpal[1], Sagar Choudhary[2], Shravan maskar[3], Shreyash Paikrao[4]

[1]*Assistant Professor, D. Y. Patil Deemed to be University, School of Humanities & Sciences, Belapur, Maharashtra, India.*

[234]*Student, Department of MCA, D. Y. Patil Deemed to be University, School of Humanities & Sciences, Belapur, Maharashtra, India.*

Abstract - *Ransomware has shifted from encrypting production data to systematically targeting the backups that organizations rely on for recovery. Cloud platforms now provide a layered set of recovery controls snapshots, versioning, immutable (WORM) storage, logically air-gapped vaults, multi-party approval, cross-account replication and isolated recovery environments but their relative merits are rarely compared on a common footing. This paper presents a structured comparative review of six cloud ransomware recovery techniques across three dimensions: data integrity, recovery performance and cost. Evidence was drawn from NIST guidance, the official documentation of AWS, Microsoft Azure and Google Cloud, industry surveys (Sophos, IBM) and peer-reviewed ransomware surveys. A twelve-criterion rubric was used to score each technique, and a simple storage-cost model was developed to show how retention and change rate drive cost. Industry data show that organizations whose backups were compromised reported median recovery costs eight times higher (USD 3 million versus USD 375,000) and were less likely to recover within a week (26% versus 46%). The comparison indicates that immutability alone is insufficient: the strongest integrity outcomes come from combining immutable storage, account-level isolation, separation of duties and clean-room validation, at the price of higher cost and operational complexity. A layered reference architecture and open research gaps are presented.*

Key Words: Ransomware recovery, Cloud backup, Immutable storage, Air-gapped vault, Data integrity, RTO/RPO, Cost analysis

1. INTRODUCTION

Ransomware remains one of the most financially damaging forms of cyber-attack. IBM's 2025 study of 600 organizations puts the average cost of an extortion or ransomware incident disclosed by the attacker at USD 5.08 million, against USD 4.44 million for the average breach of any kind [3]. Sophos' survey of 3,400 IT and security professionals reports an average recovery cost of USD 1.53 million, excluding any ransom paid [1].

Because backups are the main alternative to paying a ransom, attackers increasingly go after them. Sophos found that 94% of victims reported attempts to compromise their backups, and that when backups were compromised the median recovery cost rose to USD 3 million, compared with USD 375,000 when backups were intact [2].

Public clouds make snapshots and replicas easy to create, but they also place backup deletion inside the same identity plane as production. An attacker who obtains privileged cloud credentials can delete snapshots as easily as virtual machines. Providers have therefore introduced immutable storage, logically air-gapped vaults and multi-party approval [11]–[18]. NIST stresses that recovered data must be accurate, complete and free of malware [4], which makes data integrity, and not only availability, the central recovery requirement.

Existing surveys concentrate on detection: in one survey of more than 150 references, 72.8% of the studies addressed detection [6]. Recovery techniques are seldom compared on integrity, speed and cost together. This paper addresses three research questions:

RQ1. How do the main cloud recovery techniques differ in the integrity guarantees they give against an attacker holding privileged credentials?

RQ2. What does the available evidence say about their recovery performance (RTO and RPO)?

RQ3. How do retention, copies and operational complexity affect cost, and which trade-offs follow?

The main contributions are: (i) a taxonomy of six cloud recovery techniques; (ii) a twelve-criterion rubric and a scored comparison; (iii) a storage-cost model that exposes the effect of locked retention; and (iv) a layered reference architecture with research gaps.

2. BACKGROUND AND RELATED WORK

2.1 Ransomware and the cloud recovery problem

Under the NIST definition of cloud computing [5], customers consume compute and storage services under a shared-responsibility model: the provider secures the infrastructure, while the customer secures identities,

configuration and data. A typical cloud ransomware intrusion starts with credential misuse or an exploited vulnerability, which was the leading root cause in Sophos' 2025 survey [1]. It then proceeds through privilege escalation, discovery of backup resources, deletion or tampering with recovery points, and finally encryption of production data.

2.2 Recovery metrics

Three properties define a good recovery: the recovery time objective (RTO), the recovery point objective (RPO) and the integrity of the recovery point. The first two are classical disaster-recovery measures. The third matters because ransomware can stay dormant for weeks, so the newest backup may already be corrupted or infected. NIST SP 1800-11 describes multiple systems working together to prevent, detect, alert on and recover from events that corrupt data [4].

2.3 Related work

Razaulla et al. [6] traced the evolution and taxonomy of ransomware and found research heavily weighted toward detection. Kapoor et al. [7] proposed the Detection, Avoidance and Mitigation (DAM) framework. NIST's practice guides cover identifying and protecting assets (SP 1800-25) [8], detecting and responding (SP 1800-26) [9], and recovering (SP 1800-11) [4], while NIST IR 8374 Rev. 1 offers a ransomware risk-management profile [10]. None of these compares cloud-native recovery controls on a common rubric, which is the gap this paper addresses.

3. METHODOLOGY

3.1 Review protocol

A protocol-driven review was carried out in October 2026. Sources were limited to (a) NIST standards and guidance, (b) official documentation and engineering blogs of AWS, Microsoft Azure and Google Cloud, (c) annual industry surveys with disclosed sample sizes, and (d) peer-reviewed ransomware surveys. A source was included if it was published in 2020 or later (except [5]), addressed backup, recovery or data integrity in ransomware or destructive-attack settings, and was publicly accessible. Marketing material without technical content was excluded. News coverage was used only to corroborate figures where the primary report page could not be read in full. The final evidence base is the 20 sources in the References: 5 NIST publications, 8 provider documents, 5 industry reports or coverage of them, and 2 peer-reviewed surveys.

3.2 Techniques under review

Six techniques were identified from the sources and are defined in Table -1. They form a progression from simple copies to isolated, validated recovery.

Table -1: Cloud ransomware recovery techniques

ID	Technique	Representative implementations
T1	Native snapshots / standard backups	Provider snapshots and backups stored in the production account
T2	Versioning and soft delete	Object versioning; Azure soft delete and enhanced (always-on) soft delete [13]
T3	Immutable (WORM) backup vault	AWS Backup Vault Lock [11]; Azure immutable vault [13]; Google backup vault [15]
T4	Logically air-gapped vault with multi-party control	AWS air-gapped vault and multi-party approval [11], [12]; Azure MUA with Resource Guard [14]; Google vault in a Google-managed project [18]
T5	Cross-account / cross-region copy	AWS cross-account and cross-Region copy [12]; replication to a DR region
T6	Isolated recovery environment with validation	Analysis of restored backups in an isolated recovery environment [17]; restore testing [11]

3.3 Threat model

The attacker is assumed to hold privileged credentials inside the production account. They can delete or alter resources and backups reachable from that account and may stay dormant before triggering encryption. They cannot compromise the provider's own control plane. A technique scores well when it keeps at least one clean, restorable copy under these conditions.

3.4 Evaluation framework

Each technique is rated on twelve criteria in three dimensions (Table -2). A criterion scores 2 if fully met, 1 if partly met and 0 if not met, so each dimension index ranges from 0 to 8. For cost, a higher score means more cost-efficient. The scores are reviewer judgements based on documented capabilities, not laboratory measurements; this is discussed in Section 7.

Table -2: Evaluation criteria

ID	Criterion	What is assessed
DI1	Deletion resistance	Backups cannot be altered or deleted before retention ends
DI2	Isolation	Separation from production account and identities
DI3	Separation of duties	Multi-party approval for destructive actions
DI4	Clean-restore validation	Malware or integrity checks before restore
RP1	Restore speed	Low RTO for typical workloads
RP2	Recovery point frequency	Low RPO through frequent points
RP3	Restore granularity	File, item or volume-level recovery
RP4	Restore testing	Support for tests and orchestration
CO1	Storage overhead	Extra storage needed for copies and retention
CO2	Transfer and restore charges	Copy, egress and retrieval fees
CO3	Operational complexity	Effort to configure and run
CO4	Lock-in / misconfiguration risk	Cost of a wrong or irreversible setting

4. TECHNIQUE ANALYSIS

Table -3 compares how the three providers implement the main controls, as described in their documentation.

Table -3: Provider capabilities (from documentation)

Control	AWS	Azure	Google Cloud
Immutability	Backup Vault Lock, governance or compliance mode [11]	Immutable vault blocks loss of recovery points [13]	Backup vault: immutable and indelible [15]
Isolation	Logically air-gapped vault, account-level isolation [11]	Resource Guard, ideally in a different tenant [14]	Vault in a Google-managed project [18]
Separation of duties	Multi-party approval for air-gapped vault [12]	Multi-user authorization (MUA) [13]	IAM roles; org policy to enforce vaults [16]
Deletion safety net	Lock prevents early deletion [11]	Soft delete, 14 extra days [13]	Enforced retention period [15]
Recovery validation	Restore testing [11]	Not assessed in reviewed sources	Analysis in isolated recovery environment [17]

T1: Native snapshots. These are quick to create and restore but sit in the same identity boundary as production. A compromised administrator can delete them, which is the behaviour Sophos observed [2].

T2: Versioning and soft delete. Azure keeps deleted VM backup data for 14 additional days at no cost, and soft delete can be made always-on so it cannot be switched off [13]. This helps against accidental or early deletion, but an attacker may wait out the window or disable the feature where it is not locked.

T3: Immutable vaults. AWS Backup Vault Lock enforces a write-once, read-many model; in compliance mode the vault becomes permanently immutable [11]. Azure immutable vaults block operations that could lead to loss of recovery points [13], and Google's backup vault keeps backups unmodifiable and undeletable for the enforced retention period [15]. Immutability still lives inside one account or tenant, and a wrongly set retention in compliance mode cannot be undone.

T4: Air-gapped vaults with multi-party control. AWS air-gapped vaults are implicitly locked, so recovery points cannot be changed even with stolen source-account credentials [11], and multi-party approval allows access even when the backup account is unreachable [12]. Azure adds Resource Guard so critical operations need a second authorization [14]. Google places the vault in a Google-managed project [18].

T5: Cross-account and cross-region copy. AWS Backup supports cross-account management and cross-Region copy [12]. Separation limits the blast radius, but a copy of corrupted or encrypted data is still corrupted unless earlier versions are retained, and every extra copy adds storage and transfer cost.

T6: Isolated recovery with validation. Google recommends analysing restored backups with security

tools in an isolated recovery environment before recovering into production [17], and AWS offers restore testing [11]. This answers NIST's requirement for malware-free recovered data [4], but it adds time, compute and skills.

5. COMPARATIVE RESULTS

5.1 Data integrity

Table -5 gives the criterion scores. T1 scores zero on every integrity criterion because it offers no protection once credentials are stolen. T3 gains the full mark for deletion resistance but only partial credit for isolation and separation of duties. T4 and T6 are the only techniques reaching 6 or more of 8: T4 through isolation and multi-party control, T6 through its validation step. T5 depends on the target configuration, so its deletion resistance is partial.

Table -4: Industry evidence used in the analysis

Indicator	Value	Ref.
Median recovery cost, backups compromised vs intact	USD 3.0M vs 0.375M	[2]
Fully recovered within a week, compromised vs intact	26% vs 46%	[2]
Median ransom demand, compromised vs intact	USD 2.3M vs 1.0M	[2]
Share of ransom demanded that was paid	98% vs 82%	[2]
Average recovery cost (excl. ransom), 2024 → 2025	USD 2.73M → 1.53M	[1], [19]
Average ransom payment, 2024 → 2025	USD 2.0M → 1.0M	[1], [19]
Full recovery within a week, 2024 → 2025	35% → 53%	[19]
Average cost of ransomware incident disclosed by attacker	USD 5.08M	[3]
Average cost of any breach, 2025	USD 4.44M	[3]
Average time to recover operations after an incident	About 100 days	[20]

Table -5: Criterion scores (2 = full, 1 = partial, 0=none)

	Data integrity				Recovery performance				Cost efficiency			
Tech.	DI 1	DI 2	DI 3	DI 4	RP 1	RP 2	RP 3	RP 4	CO 1	CO 2	CO 3	CO 4
T1	0	0	0	0	2	2	1	0	2	2	2	2
T2	1	0	0	0	2	2	2	0	1	2	2	2
T3	2	1	1	0	1	2	1	1	1	2	1	0
T4	2	2	2	1	1	2	1	1	1	1	0	0
T5	1	2	1	0	1	1	1	1	0	0	1	1
T6	1	2	1	2	1	1	2	2	0	1	0	1

5.2 Recovery performance

Public documentation gives no vendor-neutral RTO figures, so performance is assessed from capabilities and from industry outcomes (Table -4). Chart -1 shows the strongest signal: only 26% of victims with compromised backups were fully recovered within a week, against 46% with intact backups [2]. Sophos' 2025 survey shows overall improvement, with 53% recovering within a week compared with 35% a year earlier [19] (Chart -2). IBM reports an average of about 100 days for organizations to recover operations after a security incident [20].

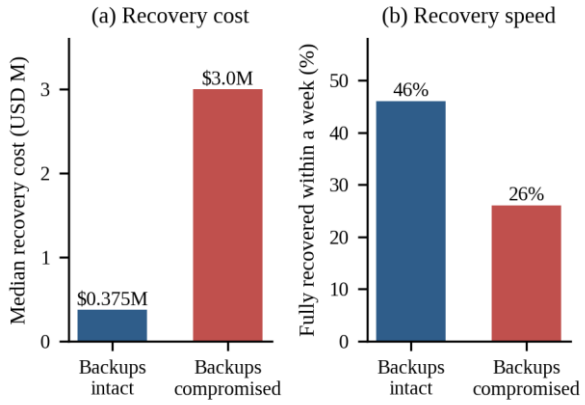


Chart -1: Effect of compromised backups on recovery (data from [2])

Recovery time is driven by five factors: the restore tier (hot or archive), data volume and parallelism, approval workflows, the time taken to validate a restore point, and whether restores have been rehearsed. T1 and T2 restore fastest because they are local. T3 and T4 add small delays from vault retrieval and approvals. T6 adds the most validation time but gives the safest cutover. Restore testing (RP4) therefore separates plans that are documented from plans that actually work.

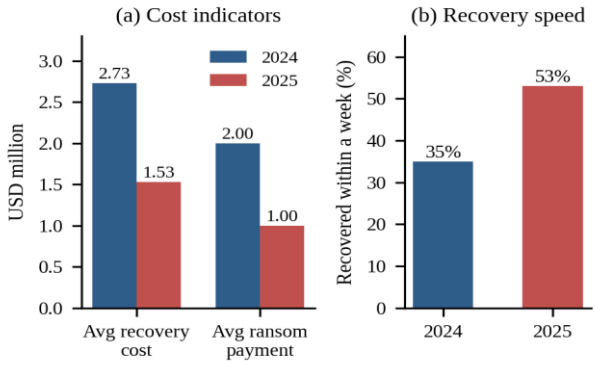


Chart -2: Year-on-year change in recovery indicators (from [1], [19])

5.3 Cost

The cost of recovery is dominated by outage duration and by whether clean backups exist. Median recovery cost was eight times higher when backups were compromised (USD 3.0 million against USD 375,000) [2], a gap of roughly USD 2.6 million. The cost of protection is smaller but depends on design. For incremental backups, storage relative to the protected data can be approximated as

$$m = n \times (1 + c \cdot R)$$

where n is the number of independent copies, c is the average daily share of data that changes and R is the retention period in days. The model ignores compression and de-duplication, so it is an upper-bound illustration rather than a price forecast. Chart -3 shows that at 10% daily change, a 30-day locked retention already needs four times the protected data, and at 90 days ten times. Because compliance-mode locks cannot be shortened [11], an over-long retention becomes a committed cost. Each extra cross-region copy (T5) multiplies m by n and adds transfer charges, while T4 and T6 add operational cost through separate accounts, approvals and recovery environments.

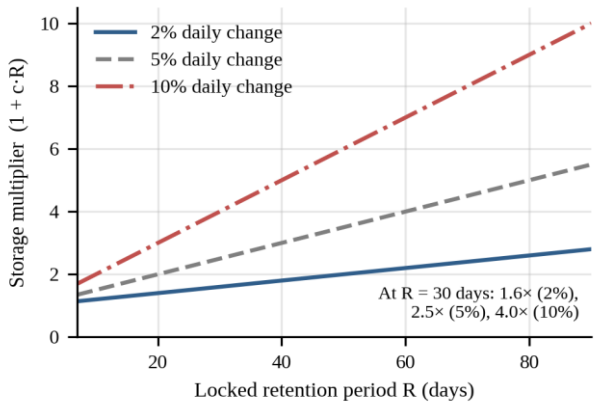


Chart -3: Storage multiplier against locked retention (model, n = 1)

5.4 Integrated comparison

Table -6 combines the indices under three weightings: integrity-first (0.5 / 0.3 / 0.2), balanced (equal weights) and cost-first (0.2 / 0.3 / 0.5). Chart -4 shows the underlying indices. Under an integrity-first view, which suits ransomware planning, T4 (5.4) and T6 (5.2) lead, followed by T3 (4.3). Under a cost-first view, T1 and T2 (both 5.5) lead, which explains why many organizations stay with simple copies until an incident occurs.

Table -6: Dimension indices and weighted scores (0–8)

Tech.	DI	RP	CO	Integrity-first	Balanced	Cost-first
T1	0	5	8	3.1	4.3	5.5
T2	1	6	7	3.7	4.7	5.5
T3	4	5	4	4.3	4.3	4.3
T4	7	5	2	5.4	4.7	3.9
T5	4	4	2	3.6	3.3	3.0
T6	6	6	2	5.2	4.7	4.0

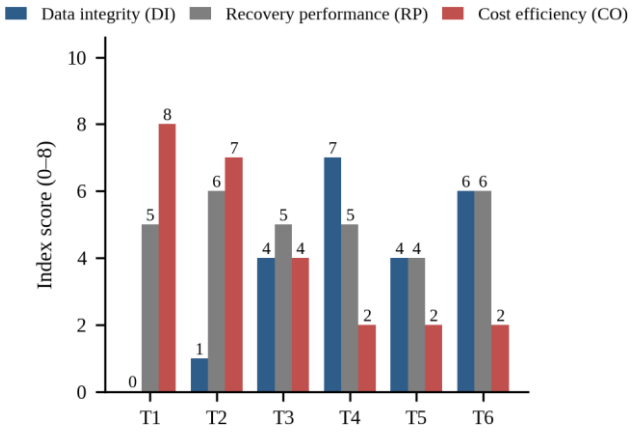


Chart -4: Data integrity, recovery performance and cost indices by technique

6. DISCUSSION

6.1 Key findings

Immutability is necessary but not sufficient. T3 stops deletion but remains within one account, so an attacker with high privileges may still reach it by other routes, and it does not tell the defender whether a restore point is clean.

Isolation and separation of duties set the leaders apart. The highest integrity scores (T4, T6) come from controls that remove the attacker's ability to act on backups at all, namely separate trust boundaries and a second approver.

Validation is the least standardized control. Only one provider's reviewed documentation describes analysing restored backups in an isolated environment [17], although NIST treats malware-free recovery as essential [4].

Cost comes from retention, copies and complexity, not from any single feature. Chart -3 shows that retention and change rate matter more than which vendor is used.

Backup integrity is economically material. An eight-fold difference in median recovery cost [2] outweighs most plausible storage costs for organizations with moderate data volumes.

6.2 Recommended layered architecture

Fig -1 shows a layered design. Frequent snapshots and versioning give low RPO and quick local restores. An immutable vault holds locked copies. A second, air-gapped copy in a separate account or region is protected by multi-party approval. Before cutover, restore points are scanned and test-restored in an isolated environment. Small organizations can start with T2 and T3 plus multi-factor authentication for backup administrators; regulated or high-value environments should add T4, T5 and T6 and rehearse restores on a schedule.

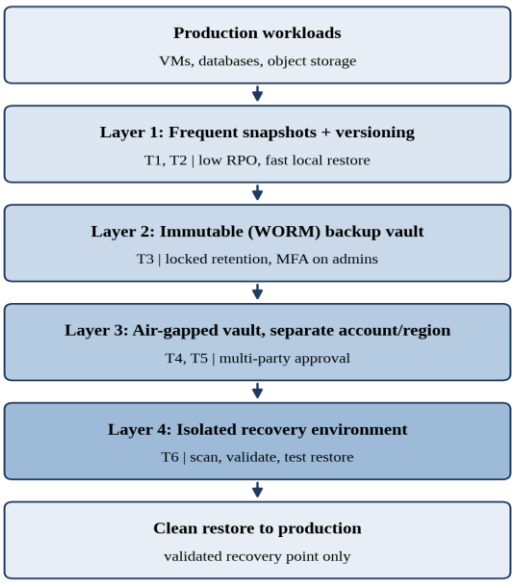


Fig -1: Layered cloud ransomware recovery architecture

7. LIMITATIONS AND FUTURE WORK

This review has five limitations. First, scores are qualitative judgements from public documentation and have not been confirmed by experiments. Second, provider documents describe intended behaviour and may favour the vendor. Third, the industry surveys rely on self-reported answers from organizations of 100 to 5,000 employees [1], and secondary coverage was used for some year-on-year figures [19], [20]. Fourth, the evidence base is limited to three providers and 20 sources, and features change quickly. Fifth, the review follows a transparent protocol but not a full database-driven PRISMA search, so a larger search across IEEE Xplore, Scopus and ACM could extend it.

Future work should (i) benchmark RTO and RPO across providers using controlled ransomware simulations, (ii) test whether validation by scanning detects dormant ransomware at scale, (iii) build cost calculators with real provider price lists, and (iv) define common metrics for backup isolation strength.

8. CONCLUSION

This paper compared six cloud ransomware recovery techniques on data integrity, recovery performance and cost. Industry evidence shows that compromised backups raise recovery cost eight-fold and slow recovery. The comparison finds that simple snapshots are cheap and fast but offer no integrity protection once credentials are stolen, that immutable vaults add real protection, and that the best results come from combining immutability, isolation, multi-party approval and clean-room validation. Cost grows mainly with locked retention, extra copies and operational complexity, so each organization should choose layers according to the value of its data and its tolerance for downtime. Vendor-neutral benchmarks are the most important gap for future research.

REFERENCES

1) Sophos, "The State of Ransomware 2025," Sophos Ltd., 2025. [Online]. Available: https://www.sophos.com/ransomware2025

2) Sophos, "The impact of compromised backups on ransomware outcomes," Sophos Ltd., 2024. [Online]. Available: https://www.sophos.com/en-us/blog/the-impact-of-compromised-backups-on-ransomware-outcomes

3) IBM Security and Ponemon Institute, "Cost of a Data Breach Report 2025," IBM Corp., Jul. 2025.

4) T. McBride, M. Ekstrom, L. Lusty, J. Sexton and A. Townsend, "Data Integrity: Recovering from Ransomware and Other Destructive Events," NIST SP 1800-11, 2020, doi: 10.6028/NIST.SP.1800-11.

5) P. Mell and T. Grance, "The NIST Definition of Cloud Computing," NIST SP 800-145, 2011.

6) S. Razaulla et al., "The age of ransomware: A survey on the evolution, taxonomy, and research directions," IEEE Access, vol. 11, 2023, doi: 10.1109/ACCESS.2023.3268535.

7) A. Kapoor et al., "Ransomware detection, avoidance, and mitigation scheme: A review and future directions," Sustainability, vol. 14, no. 1, art. 8, 2022, doi: 10.3390/su14010008.

8) National Institute of Standards and Technology, "Data Integrity: Identifying and Protecting Assets Against Ransomware and Other Destructive Events," NIST SP 1800-25.

9) National Institute of Standards and Technology, "Data Integrity: Detecting and Responding to Ransomware and Other Destructive Events," NIST SP 1800-26.

10) National Institute of Standards and Technology, "Ransomware Risk Management: A Cybersecurity Framework 2.0 Community Profile," NIST IR 8374 Rev. 1.

11) Amazon Web Services, "Securing backup data against modern threats with AWS Backup," AWS Storage Blog. [Online]. Available: https://aws.amazon.com/blogs/storage/securing-backup-data-against-modern-threats-with-aws-backup/

12) Amazon Web Services, "AWS Backup features." [Online]. Available: https://aws.amazon.com/backup/features

13) Microsoft, "Overview of security features in Azure Backup," Microsoft Learn. [Online]. Available: https://learn.microsoft.com/azure/backup/security-overview

14) Microsoft, "Azure Backup data protection best practices," Microsoft Docs. [Online]. Available: https://docs.azure.cn/en-us/backup/azure-backup-data-protection-best-practices

15) Google Cloud, "Backup vault for immutable and indelible backups," Backup and DR Service documentation. [Online]. Available: https://docs.cloud.google.com/backup-disaster-recovery/docs/concepts/backup-vault

16) Google Cloud, "Best practices for ransomware protection with Backup and DR," Backup and DR Service documentation. [Online]. Available: https://docs.cloud.google.com/backup-disaster-recovery/docs/concepts/best-practices-ransomware

17) Google Cloud, "Backup and DR Service," product documentation. [Online]. Available: https://cloud.google.com/backup-disaster-recovery

18) Google Cloud Blog, "Backup vaults add support for disk backup and multi-region," Jun. 17, 2025. [Online]. Available: https://cloud.google.com/blog/products/storage-data-transfer/backup-vaults-add-support-for-disk-backup-and-multi-region

19) SecurityBrief UK, "UK firms face rising ransomware costs but recover faster," 2025 (coverage of Sophos State of Ransomware 2025). [Online]. Available: https://securitybrief.co.uk/story/uk-firms-face-rising-ransomware-costs-but-recover-faster

20) CSO Online, "What is the cost of a data breach?" (coverage of IBM Cost of a Data Breach Report 2025). [Online]. Available: https://csoonline.com/article/567697/what-is-the-cost-of-a-data-breach-3.html