WHITE PAPER

Proven Recovery
Three tiers. One chain of evidence.

Proven Recovery: Independent by Architecture

No single supplier can make your recovery independent. Here's how Bacula, Nexsan,and Spectra Logic build it together.

When an incident takes out production, a backup is only as good as everything it depends on: the software, the catalog, the credentials, the storage, and the supplier behind each one. Concentrate those in one vendor's stack and a single failure can reach every copy.

This paper sets out a recovery architecture built so that can't happen. Nexsan Unity provides a fast protection repository and a retention-locked compliance copy with S3 Object Locking. Spectra Logic tape provides an isolated vault whose media leaves the network. Bacula Enterprise drives all three, and can be rebuilt from any of them. The paper explains why each layer exists, the order recovery has to follow, how to prove it works, and what the architecture does not solve.



EXECUTIVE SUMMARY

Every backup architecture contains an assumption that is rarely written down: that when recovery is needed, the things recovery depends on will still be there. The software vendor will be reachable. The storage will respond. The credentials will work. The catalog will be readable. Most of the time that assumption holds, which is why it survives unexamined.

The incidents of the last two years have been examinations. A cloud provider told customers it was unable to restore access to data held exclusively in a damaged region.[1] A national computer emergency response team recorded a destructive attack in which the backup and restore machinery was itself the delivery route for the wiper.[2] An operator lost three data centers in four days and faced equipment replacement measured in months.[3] In each case the organizations that continued were the ones holding a copy that was out of reach of the incident.

This paper makes an architectural argument rather than a product one. A recovery capability assembled from a single supplier's stack inherits that supplier's failure modes: one software defect class, one credential boundary, one commercial relationship, one jurisdiction. Independence is not a feature that can be added to such a stack. It is a property of how the layers are separated, and it has to be designed in.

The argument runs in four steps, each developed below:

- **Immutability is not isolation.** A retention-locked copy is protected from alteration. It is not out of reach. The governing definition of an air gap concerns connection, not permission.[4]
- **A copy is not a capability.** If reading the copy depends on the catalog and configuration that lived in the facility you are planning to lose, you hold data you may not be able to address inside the time you have committed to.
- **One supplier cannot supply independence from itself.** Whatever the contract says, a single vendor's stack has a single set of ways to go wrong.
- **Only a test settles it.** A recovery capability that has been asserted and one that has been demonstrated look identical on an architecture diagram and behave differently on the day.

What follows describes an architecture built from three storage layers under separate ownership: a protection repository, a retention-locked compliance copy and a tape vault whose media leaves the network. These layers are driven by one protection system that can be rebuilt from any of them. It also sets out, in section 9, what this architecture does not solve.

01 The assumption inside every backup architecture

Ask an operator whether they can recover and the answer is almost always yes. Ask what recovery depends on and the conversation changes, because the dependencies are rarely enumerated. They typically include: the backup software being installed and licensed; the catalog that maps logical objects to physical media; the configuration describing clients, schedules and storage; the credentials for the storage; the encryption keys; and the hardware or cloud capacity to restore onto.

Each of those is a dependency on something that may not survive the event. The question that separates a recovery plan from a recovery capability is not whether a copy exists. It is which of those dependencies the copy carries with it, and which it assumes.

THE DISTINCTION THIS PAPER TURNS ON

An offsite copy answers the question “is the data somewhere else?” A recovery capability answers the question “can this data be read, by us, without the environment that produced it, and how long does that take?”

02 Three failure modes a single stack cannot separate

A protection architecture sourced entirely from one supplier is efficient to buy and support. It also concentrates three distinct categories of failure into one place.

Failure mode	How it presents	Why one supplier cannot separate it
Technical	A software defect, a firmware fault, or a design limitation that affects how data is written or read	The same code and the same formats are in every layer. A defect that makes one copy unreadable is likely to affect the others
Operational	A credential, an administrative console or an automation path that reaches further than intended, including the backup system itself being used as a route	A single management plane is a single boundary. Immutability constrains what can be done to a copy; it does not change what can reach it
Commercial and jurisdictional	Acquisition, end of support, a licensing dispute, an export restriction, or a legal order compelling disclosure or denying service	One relationship, one corporate parent, one jurisdiction. The contract is with the same party across every layer

None of these is a criticism of any particular supplier, ours included. They are structural properties of concentration. The only remedy for concentration is separation.

03 What isolation actually requires

The term air gap is used loosely and it has a governing definition. **CNSSI 4009, carried in the NIST glossary, defines it as an interface between two systems at which they are not connected physically and any logical connection is not automated.** Two clauses, about a single interface. Not an architecture, not a facility, not a security program. Where a claim is tested, in an assessment, in a contract dispute, after an incident, that is the text it is tested against.

Read against that definition, three things commonly sold as isolation are not:

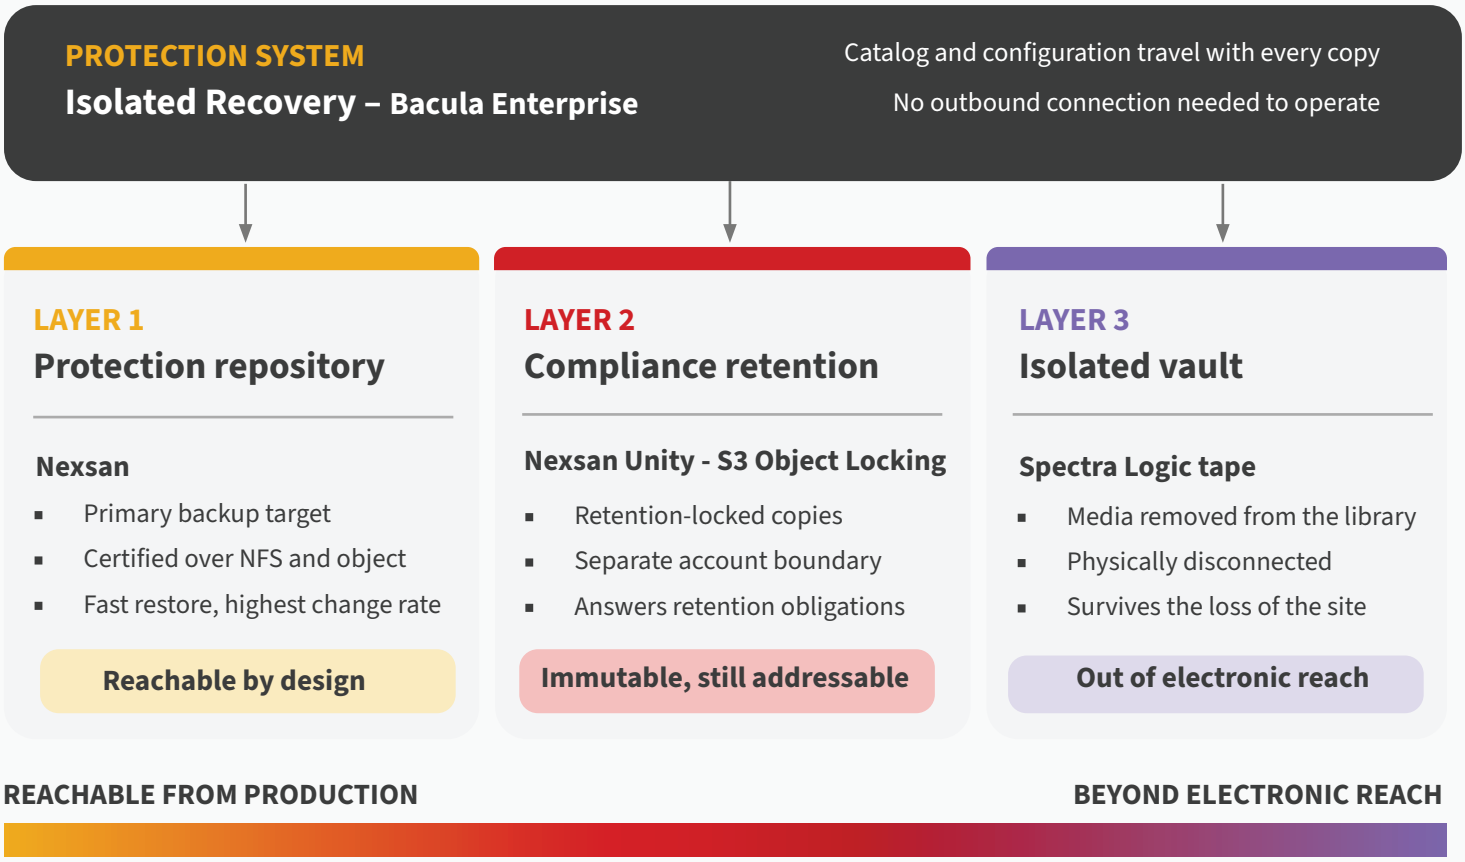
- **Immutable object storage that production can address.** Immutability protects a copy from alteration. The account that can write it is an automated logical connection.
- **A replica in a second region of the same provider.** Geographic separation is not isolation when the control plane, the identity provider and the credentials are shared.
- **A backup system reachable from the production domain.** If a privileged directory account can reach the backup server, the backup server is inside the blast radius whatever the storage behind it does. The December 2025 wiper case turned on exactly this.[2]

This is not an argument that immutability is worthless. It is an argument that immutability and isolation answer different questions, and that an architecture needs both answers.

04 Why three layers

If concentration is the problem and separation is the remedy, the question becomes how many layers and where to draw the lines. Three is not arbitrary. It is the smallest number that gives each of the three failure modes in section 2 somewhere different to fail.

One protection system. Three storage layers, three failure modes.



Each layer fails differently. That is the point: no single event, credential or supplier reaches all three.

The layers are ordered by reachability rather than by cost or speed. The repository is reachable by design, because the ordinary incident, such as a deleted file, a corrupted database, or a failed host, needs a fast path. The retention-locked copy sits behind a separate account boundary and resists alteration. The vault is written to media that is then removed, which is the only one of the three that satisfies the definition in section 3.

Each layer is bought from its own supplier, and that is the point rather than a procurement inconvenience. Separate vendors mean separate code, separate support relationships, separate commercial exposure and separate administrative boundaries. The architecture is only as independent as its most concentrated layer.

05 The layers, and what each one answers

Layer and supplier	The question it answers	What it does not answer
Protection repository **Nexsan**	Can we restore quickly from the copy we use most? Certified with Bacula Enterprise over both NFS and object	Nothing about isolation. It is reachable, because a fast path has to be
Compliance retention **Nexsan Unity with S3 Object Locking**	Can a copy be held unchanged for a stated period, under a separate account boundary, to meet a retention obligation?	Whether the copy is out of reach. Locked is not unreachable
Isolated vault **Spectra Logic tape**	Does a copy exist that no credential, process or service in the production environment can address? Media removed from the library is physically disconnected	Speed. This is the copy that exists so that the other two can both be wrong
Protection system **Bacula Enterprise with Isolated Recovery**	Can the thing that reads the copies be rebuilt from the copies themselves? The catalog and configuration travel with the data, and operation needs no outbound connection, no license call-home, no entitlement check	Placement, media handling and credential separation. No software makes a copy isolated

06 The recovery sequence

The order of operations is the part most designs get wrong, and it is invisible on an architecture diagram.

The test that defines it - recovery with Bacula Systems treated as unavailable

1. **Isolate** — Confirm the copy is unreachable from production
2. **Verify** — Establish integrity and readability of the media
3. **Rebuild** — Protection system first, from the media, keys from your custody
4. **Validate** — Services and applications checked in isolation
5. **Record** — Timed and dated, recovery time and recovery point apart

The first thing recovered is not data. It is the thing that knows how to read the data, the protection system, its configuration and its catalog. An organization whose index and configuration lived only in the destroyed facility holds a copy of the data and not yet a recovery capability.

That position is recoverable rather than fatal. Catalog records can be reconstructed from the volumes themselves, given the right setup, the media and the keys. But reconstruction is slower and less certain than reading a catalog that survived, and it may exceed the recovery time the organization has committed to. The risk to plan against is a slow or uncertain recovery, not a guaranteed loss.

07 Proving it

Everything above is a claim until it has been exercised. The test that settles it has a specific shape: production is treated as gone, and **every supplier in the architecture is treated as unreachable, including the software vendor.**

The customer's own team recovers an agreed set of systems using only retained software, their own documentation, their own keys and an isolated copy. It is timed and witnessed. The exercise runs the complete startup path, including catalog recovery, workload startup, application validation, and records recovery time and the recovery point actually achieved separately, because restored files or booted virtual machines on their own demonstrate neither that a business service is back nor that the restored state is clean.

The result is a dated record, and the record is the deliverable. It is not regulatory evidence in itself; it is the raw material a compliance process, an auditor or a board can use. Most organizations find something on the first run, which is the argument for running it before an incident rather than during one.

08 Regulatory context, scoped honestly

Recovery obligations have tightened across several jurisdictions, and the direction of travel is consistent: segregate the copy, test the restoration, state the time objective. Vendor material routinely overstates the specifics, and a buyer is better served by the accurate version.

Instrument	What it actually requires	On whom
DORA, Art. 12(3)[5]	Restoration using ICT systems physically and logically segregated from the source system. A particularly explicit EU requirement on segregated restoration, of the systems used to restore, not a duty to keep every backup offline.	Financial entities. Directly applicable across the EU since 17 January 2025
NIS2, Art. 21(2)(c)[6]	Risk-based business continuity measures including backup management and disaster recovery, and assessment of their effectiveness. The detailed separation and testing rules sit in an implementing regulation with a narrower scope.	Essential and important entities, through national transposition
NERC CIP-009-6[7]	Test recovery plans every 15 calendar months; test a representative sample of backup information for usability every 15 months; full operational exercise every 36 months.	US bulk electric system. Enforceable with penalties, and has been for a decade
NCA ECC-2:2024, 2-9[8]	Backup scope covering critical assets, ability to recover quickly after a cybersecurity incident, and periodic testing of recovery effectiveness. It does not mandate offline or isolated storage and sets no test frequency.	Saudi government agencies and private operators of critical national infrastructure

09 What this architecture does not solve

- **It does not prevent an incident.** Nothing here reduces the likelihood of an intrusion, an outage or a physical event. It changes what the consequence is.
- **It does not remove the organizational requirement.** A named owner for the vault, a written procedure, staff who have performed the recovery rather than read about it, and a schedule that puts them through it are not properties a product can supply. A verification exercise that restores files while the production backup server quietly does the work has tested nothing about isolation.
- **It does not produce regulatory evidence by itself.** It produces recovery evidence that a compliance process can use. Bacula Systems holds no ISO/IEC 27001, ISO 9001, SOC 2 or TISAX certification and no independent penetration test report; where a tender requires third-party attestation we say so.
- **It does not make the slow layer fast.** Tape that has left the library takes time to retrieve and read. That is the cost of the property it provides, and it is why the architecture has three layers rather than one.

None of these limits is an argument against the architecture. They are the terms on which it works, and an operator who knows them will build something that survives contact.

10 References

Cited in order of first appearance.

1. The Register, “AWS says wartime damage means some Middle East cloud resources are gone for good,” 16 September 2026, reporting AWS Health Dashboard statements; and Help Net Security, “AWS Middle East outage: permanent data loss in Bahrain and UAE,” 17 September 2026.
2. CERT Polska, Energy Sector Incident Report — 29 December 2025. The attack script backed up the default domain policy, modified it to add a privileged scheduled task running a wiper from a network share, and then restored the tampered backup in order to activate it.
3. The Register, “Ukrainian drone strikes on Yandex datacenters knock cloud and payments offline,” 12 October 2026; and Meduza, reporting Kommersant on equipment replacement of up to six months, 12 October 2026.
4. Committee on National Security Systems, CNSSI No. 4009, Committee on National Security Systems Glossary, definition of “air gap,” adopted from IETF RFC 4949 and carried in the NIST glossary.
5. Regulation (EU) 2022/2554 (DORA), Article 12(3); Article 11(6)(a) on testing at least yearly, including switchover to backups. Applicable from 17 January 2025.
6. Directive (EU) 2022/2555 (NIS2), Article 21(2)(c) and (f); Commission Implementing Regulation (EU) 2024/2690 of 17 October 2024, Annex points 4.1 and 4.2, whose scope is limited to specified digital and ICT service providers.
7. North American Electric Reliability Corporation, CIP-009-6, Cyber Security — Recovery Plans for BES Cyber Systems, Requirements R1 and R2.
8. National Cybersecurity Authority (Saudi Arabia), Essential Cybersecurity Controls ECC-2:2024, Subdomain 2-9, Backup and Recovery Management, controls 2-9-1 to 2-9-4. Offline and offsite backup appear in Appendix A as definitions rather than as requirements.

Incidents are cited to illustrate failure and recovery patterns, not to characterize any conflict or any party to one; this paper makes no attribution. Reporting on active conflicts is incomplete and in places disputed. This document is not legal advice, and the regulatory summaries in section 8 are not a substitute for advice on any organization’s own obligations. Nothing in it constitutes a quotation or an offer.

DELIVERED BY

NEXSAN | SPECTRA | Bacula Systems