SCALITY

THREE PRODUCTS · ONE FOUNDATION · TWO FLAGSHIP EXTENSIONS

# One proven object-storage core, shaped for *autonomous AI*, *cyber resilience* and *sovereign cloud*.

**Scality ARTESCA protects data. Scality RING scales it. Scality ADI orchestrates it.** The same distributed object-storage engine takes three shapes for three different problems, so you standardize on one data architecture and one operational model, then deploy the shape each workload needs. Two flagship extensions, ARTESCA+ and RING XP, cover specialized backup and AI-performance cases.

## THE PORTFOLIO AT A GLANCE

BACKUP

### ARTESCA

**Built for backup**

A simple, secure S3 target for immutable data protection. A software appliance with an embedded hardened OS: deploy in minutes, scale to petabytes.

SOVEREIGN CLOUD

### RING

**Built for sovereign cloud at scale**

Multi-petabyte to exabyte unified file and object storage. Pure software-defined, bring your own Linux. Proven in MSP sovereign storage, data lakes and cloud-scale platforms.

ENTERPRISE AI · LATEST

### ADI

**built for AI**

**AI Data Infrastructure, a software** appliance with embedded ScalityOS. Four optional temperature tiers and every protocol, from GPU-direct flash to deep cold, in one namespace.

FLAGSHIP EXTENSION

### ARTESCA+

Veeam and ARTESCA delivered as a single appliance, so the immutable S3 target is built in rather than bolted on. No separate Veeam server. Available in 3 or 6 node architectures, optimized for protecting 20 TB to 440 TB usable.

FLAGSHIP EXTENSION

### RING XP

All-NVMe AI tier. Under 511 µs GET and under 741 µs PUT for the majority of 4 KB object requests, which Scality states is 10–20× faster than Amazon S3 Express One Zone and 50–100× faster than standard Amazon S3.

Source: Scality, 2024-10-30, RING / RING XP 9.3.0.8 test environment.

BUILT IN · NO ADD-ON

### Guardian + MCP

Scality Guardian operational intelligence: read-only by design, surfacing recommendations and real-time power awareness at system, node and workload level. MCP-enabled extensibility lets you drive ADI operations from your own AI stack. Included with ADI on embedded ScalityOS.

## SHARED ACROSS THE FAMILY

**Native S3 API**
Drop-in AWS S3 compatibility. Apps, SDKs and policies work unchanged.

**CORE5 cyber resilience**
Five layers of protection against ransomware: API, data, storage, geographic and architectural.

**Scale without downtime**
Add nodes or drives online. No migration, no re-architecting.

**Single namespace**
One view across sites and clouds, with more than 6 trillion objects under management in production.

**Data services built in**
Versioning, Object Lock, lifecycle and replication. Native, not bolted on.

## DEPLOYMENT AND COMMERCIAL MODEL

| ARTESCA | RING | ADI |
|---|---|---|
| › Use: immutable backup target and ransomware-proof restore | › Use: cloud-scale data lakes, MSP sovereign storage and content repositories | › Use: enterprise AI, cyber resilience and sovereign control |
| › Performance: scales with node count; per-node throughput is published in the ARTESCA documentation | › Scale: multi-petabyte to exabyte in one logical namespace | › Scale: multi-petabyte to exabyte across tiers |
| › Scale: 20 TB to 8.5 PB, single, three or six node | › Performance: 420 GB/s sustained read and 250 GB/s sustained write on HDD across just over 100 nodes, concurrent file and object | › Tiers: four optional, GPU-Direct (sub-50 µs) › HOT (multi-TB/s) › WARM › COLD |
| › Media: HDD or SSD, embedded OS | › Resilience: multi-site geo, cross-region replication | › Protocols: S3, NFS, SMB, GPU-direct |
| › Ecosystem: Veeam, Commvault, Cohesity, Rubrik, HYCU, Kasten, Veritas | › Model: pure software-defined, customer-managed Linux, standard x86 (Intel or AMD) | › Operations: Guardian recommendations and MCP extensibility, human-in-the-loop |
| › Assurance: backed by a $100,000 Cyber Guarantee* | | |
| › Software appliance, embedded hardened OS | › Pure software-defined, customer-managed Linux | › Software appliance, embedded ScalityOS, minimum 3 servers |
| › Single, three or six node, deploy in minutes | › Standard x86 (Intel or AMD), your hardware, your control | › On-premise or third-party datacenter |
| › Dark-site ready, offline install, no Internet | › Dark-site ready, offline install, no Internet | › Dark-site ready, offline install, no Internet |
| › Annual subscription | › Capacity-tier subscription | › Outcome-based SLA or capacity-tier subscription |

© 2026 Scality

**Built for GPU hunger, deep cold and everything in between.**

## 01 Standardize once, deploy per problem

Because all three share one storage engine and the S3 API, your teams learn one model, your data protection behaves the same way everywhere, and data moves between shapes without re-platforming. You consolidate vendors without giving up fit-for-purpose deployment.

Proof: more than 12 EB in production on that one engine, across 70+ countries.

## 02 Temperature, not averages

ADI aligns media and architecture to each workload's heat, GPU-direct flash for AI and deep cold for archive, and policy moves data down the gradient automatically. You stop buying peak performance for cold data, so TCO is optimized per workload rather than averaged across the estate.

Proof: four optional tiers, one namespace. Up to 20× better power efficiency than all-flash for typical enterprise data distributions.

## 03 Immutability is in the engine

The storage engine never overwrites in place. Every write lands at a new address, so ransomware-proofing is not a policy you can misconfigure away, it is how the object store works. CORE5 cyber resilience across API, data, storage, geographic and architectural layers is shared DNA across all three products.

Proof: append-only engine, CORE5, Object Lock compliance mode.

## 04 Operations that do not scale with capacity

As estates grow into AI, the operational burden rather than the hardware becomes the ceiling. Guardian surfaces recommendations and real-time power awareness at system, node and workload level, and MCP-enabled extensibility lets you integrate your own AI tools and automation workflows directly into ADI operations, so the platform can be driven by your own AI stack. Every operation follows a human-in-the-loop principle.

Proof: Guardian and MCP-enabled extensibility both shipped with ADI at general availability on 2026-05-12, included on embedded ScalityOS with no add-on.

### CHOOSING WITHIN THE FAMILY

### Choose ARTESCA

- You need a turnkey, immutable S3 backup target
- You run Veeam, Commvault, Cohesity or Rubrik
- Mid-market scale, 20 TB to 8.5 PB
- Immutability is the number-one requirement

### Choose RING

- You want to manage your own OS on standard x86 (Intel or AMD)
- Your team already runs Linux at scale
- Cloud-scale file and object on shared infrastructure
- RING XP when AI training needs sub-millisecond flash

### Choose ADI

- AI pipelines need GPU-direct and cold archive together
- Power, cooling or flash supply is a constraint
- Cyber, sovereignty and AI must be one architecture
- You prefer an appliance with built-in operational intelligence

**>12 EB**
IN PRODUCTION

**1000+**
ENTERPRISES IN PRODUCTION

**Since 2009**
CONTINUOUS PRODUCTION

**ISO 27001**
CERTIFIED

*The $100,000 Cyber Guarantee is paid directly to the customer if an external cyberattack destroys or encrypts data stored immutably on ARTESCA. All of the following conditions are required: a commercial licence in production use, a 50 TB minimum licence, ARTESCA v4.1.3 or higher kept current with the latest release, protected data stored using S3 Object Lock in compliance mode, and recommended configuration and deployment practices followed. Specifications reflect the current product family as of 2026-09-11; Scality ADI reached general availability on 2026-05-12. Veeam, Commvault, Cohesity, Rubrik, HYCU, Kasten and Veritas are technology and alliance partners. Dark-site ready means all three install and update without Internet access, using an offline installer and a local repository; a dark site is an environment with no Internet access. RING XP latency and Amazon S3 comparison figures are Scality-published, dated 2024-10-30, measured on RING / RING XP 9.3.0.8.

© 2026 Scality