

Roadrunner SMB™

Windows SMB File Shares on Amazon EFS — Without Windows File Servers

Architecture & Design Whitepaper

Version **1.3** · 24 August 2026

Copyright © 2026 Roadrunner SMB, LLC. All rights reserved.

Roadrunner SMB and Elastic SMB are trademarks of Roadrunner SMB, LLC.

1. Executive Summary

Windows file servers remain one of the most operationally complex workloads in enterprise cloud migrations. **Roadrunner SMB (RRSMB)** delivers **Elastic SMB for Amazon EFS** — native **SMB3** file services on **Amazon Elastic File System (EFS)** for teams who already run elastic, AWS-native storage and need Windows file access without operating Windows file servers or hand-building Samba clusters.

Customers subscribe in AWS Marketplace, launch a **CloudFormation** template into their VPC, and manage shares from a browser-based **Admin UI**. Version **1.3** has two production outcomes: **login-scale hardening** (8 NFS mux channels per share, accept-path admission, FSLogix login characterization) and a **stronger administrative security posture** (Eval Admin behind CloudFront + ALB + WAF; Dual-AZ Production Admin not internet-facing).

RRSMB is particularly well suited for:

- **VDI / FSLogix** login and profile storage on AWS (characterized on real Windows Server clients running FSLogix; WorkSpaces sizing is inferred for planning)
- **Departmental file shares** and home directories
- **MSP-managed** AWS file services
- **Windows file server modernization** with Active Directory and Windows ACL semantics preserved

Roadrunner SMB is designed around **performance, stability, flexibility, active-active scale-out, and elastic Amazon EFS storage** with automatic Infrequent Access lifecycle for aging files. The durable cost distinction is **Amazon EFS Elastic Throughput**: login-storm demand is paid as it happens, rather than a 24x7 provisioned throughput tier sized for the Monday-morning peak. What that is worth depends on the deployment, so this document states the mechanism rather than a savings percentage.

Deployment model: RRSMB runs entirely in the **customer's AWS account** using standard AWS services — Amazon EC2, Amazon ECS, Amazon EFS, Network Load Balancing, Amazon DynamoDB, and AWS Secrets Manager. No proprietary filesystem, external control plane, or customer data egress is required.

Pricing: RRSMB software is billed through **AWS Marketplace as hourly consumption** — active Windows **ShareHours** and mounted EFS **ManagedStorageGBHours**. Current rates are published on the [Roadrunner SMB website](#) and may change independently of this document. AWS infrastructure (EC2, EFS, NAT, load balancing, data transfer, and related services) is billed separately by AWS.

RRSMB is **not** intended to replace every Windows file server scenario. Workloads requiring native Windows features such as VSS/shadow copies, DFS Replication, database storage over SMB, SMB Multichannel, or deep Windows kernel integration may be better served by alternatives such as **Amazon FSx for Windows File Server**.

2. The Problem Space

Organizations migrating Windows workloads to AWS frequently require SMB-based file services for home directories, collaboration shares, application storage, and VDI profile containers. Common patterns — self-managed Windows file servers on EC2, managed Windows file services, and third-party NAS appliances — typically rely on always-on server infrastructure, fixed throughput allocation, and manual high-availability configuration.

Metadata-heavy workloads amplify the challenge. SMB is chatty under high concurrency (for example FSLogix login storms). Small-file operations generate frequent metadata lookups and ACL evaluations. When underlying infrastructure is not optimized for these patterns, perceived performance degrades disproportionately.

Amazon EFS provides durable, elastic, regionally resilient storage with lifecycle tiering — but exposes **POSIX/NFS** semantics. Windows workloads requiring SMB cannot consume EFS directly without an intermediary SMB layer.

Roadrunner SMB closes this gap by combining **clustered Samba** with **EFS as the authoritative storage substrate**, **Active Directory** integration, and **AWS-native** deployment and operations.

2.1 Comparative context: RRSMB and Amazon FSx for Windows File Server

The following comparison is intended to help architects understand where each solution fits — not to position one as universally superior to the other.

Dimension	Roadrunner SMB	FSx for Windows File Server
Underlying storage	Amazon EFS (elastic, lifecycle tiering)	Dedicated Windows file system (NTFS)
SMB layer	Clustered Samba on Amazon Linux 2023	Native Windows Server
Data sovereignty	All data and control plane in customer AWS account	AWS-managed service
Multi-protocol (NFS + SMB)	Because files live as standard files in Amazon EFS, Linux and Unix clients can access the same data over NFS while Windows users connect over SMB — one dataset, no second storage system. Ideal for build pipelines and data-processing estates alongside Windows shares. (Note: VDI/FSLogix profile workloads should remain SMB-only.)	SMB native; NFS requires separate architecture
Cluster sizing (1.3)	Dual-AZ production 2-8 nodes (default 2); 9-16 Experimental ; Eval Kit is evaluation only	AWS-managed capacity model
Active Directory	Customer-managed AD (on-prem, EC2, or AWS Managed AD)	AWS Managed AD or self-managed AD
VSS / shadow copies	Not supported	Supported
DFS Replication	Not supported	Supported

Dimension	Roadrunner SMB	FSx for Windows File Server
RRSMB software billing (GA)	Pay for hourly consumption (ShareHours + ManagedStorageGBHours); current rates on roadrunnersmb.com	FSx pricing (see AWS)

Roadrunner SMB is well suited when you need EFS-backed elasticity, the same EFS dataset over NFS and SMB (separate protocol access), full in-account operational transparency, and pay-as-you-go hourly consumption on AWS Marketplace.

FSx for Windows File Server is often a better fit when you require native Windows features (VSS, DFSR) and prefer a fully AWS-managed Windows file system.

3. Design Principles

3.1 AWS Native-First

RRSMB is built from standard AWS primitives and deployed via **CloudFormation** into the customer's VPC through **AWS Marketplace**.

3.2 Customer data sovereignty

All file data, ACL metadata (in customer-owned DynamoDB), and configuration remain in the **customer AWS account**. RRSMB does not operate an external control plane or exfiltrate customer file contents.

3.3 Replaceable compute, durable storage

RRSMB nodes are **replaceable ECS tasks**. Customer file data resides on **customer-owned EFS**. Cluster coordination uses DynamoDB and EFS-backed CTDB artifacts. Node replacement and rolling upgrades do not require storage migration.

3.4 Consistent SMB semantics

Samba is authoritative for SMB protocol behavior and **NT ACL** enforcement. **Active Directory** provides authentication (Kerberos/NTLM). Windows clients see familiar share semantics without running Windows Server for the SMB gateway layer.

3.5 Control plane vs. data plane separation

Administrative traffic (Admin UI, health, support reports) is logically and network-separated from **SMB data** traffic. Security groups and listener CIDR parameters enforce the boundary.

3.6 Elastic SMB metadata acceleration

RRSMB includes **Elastic SMB™** metadata acceleration (patent pending) to reduce EFS metadata latency for profile-heavy and login-storm workloads — improving VDI responsiveness without changing Windows client behavior. Because caching absorbs login-storm load, deployments avoid pre-provisioning the peak throughput a Windows file server requires, paying instead for actual consumption.

3.7 Observability by design

Integrated metrics, structured logging, and **Support Report** generation from the Admin UI simplify day-2 operations and vendor support.

4. Architectural overview (1.3 Marketplace)

RRSMB deploys into the customer's **existing VPC** via Marketplace CloudFormation. The stack performs **VPC discovery and validation** before creating product resources. Unsuitable VPCs fail fast with **VPC Unsuitable: ...** and roll back.

4.1 Buyer flow

1. **Subscribe** in AWS Marketplace
2. **Launch** the CloudFormation template (VPC ID + node count)
3. **Complete First-Time Setup** in the Admin UI (domain join, TLS, first share)

4. Mount shares from Windows clients on allowed networks

Most teams reach a first SMB share in about **30 minutes** after stack creation.

4.2 High-level block diagram




Figure 4-1 · High-level component view

4.3 Node exploded view

Each RRSMB node is an **ECS task** on Amazon Linux 2023 containing:

- **Samba / VFS** — SMB3, NT ACL enforcement, Kerberos/NTLM
- **CTDB** — clustered Samba state coordination
- **winbind** — Active Directory integration
- **Management server & Admin UI backend**
- **Elastic SMB** metadata acceleration layer
- **Metrics** sidecar for health and observability




Figure 4-2 · RRSMB node — ECS task exploded view

4.4 Eval Kit vs Dual-AZ

	Eval Kit	Dual-AZ (production)
Purpose	Technical evaluation / POC only	Only supported production path
VPC / AD	Kit-created	Customer VPC and AD
Nodes	1-8 supported (default 2); 9-16 Experimental	1-8 supported (default 2); 9-16 Experimental
Admin	Internet-accessible via CloudFront + ALB + WAF (no public Admin NLB)	Not internet-facing ; private ALB + WAF only
SMB	Private NLB :445	Private NLB :445
Secure Admin path	Not required (evaluation convenience)	Customer VPC access model: VPN, Direct Connect, jump host, or equivalent

Single-AZ production Marketplace launchers are withdrawn.

4.5 Networking model (1.3)

Hardened Admin security posture. The Eval Kit places the Admin UI and API behind CloudFront, ALB, and AWS WAF rather than exposing Admin on a public NLB. Dual-AZ Production keeps the Admin UI and API private, with access limited to the customer's VPC/private network. Customers and MSPs use their existing approved network-access model (VPN, Direct Connect, jump host, or equivalent). Roadrunner inherits the customer's VPC security posture for production administration.

Eval Kit: internet-accessible Admin is permitted, but is protected by CloudFront + ALB + WAF. Do not publish the ALB DNS. Use stack output `AdminUIUrl` .

Dual-AZ: Admin UI/API are not internet-facing. RRSMB provides the Admin endpoint on the private network. The customer provides the secure path.

SMB service traffic remains private behind an NLB in both deployment tiers. SMB clients connect from VPC, VPN, Direct Connect, or other ranges allowed in `SmbClientCidr` .

RRSMB **nodes and EFS** remain in **private subnets**. EFS file systems created by the Eval Kit and Dual-AZ templates default to **Elastic** throughput; Bursting remains an override. Each customer share uses **8** NFS / efs-proxy channels (`NfsChannels`). The supported setting is `NfsChannels=8` ; do not reduce it.




Figure 4-3 · Eval Kit — CloudFront + WAF Admin; private SMB NLB




Figure 4-4 · Dual-AZ — private Admin ALB/WAF; private SMB NLB; customer-provided access path

4.6 Control plane vs. data plane

- **Data plane:** Windows / FSLogix SMB clients → private NLB TCP/445 → Samba → EFS (8-channel NFS mux)
- **Control plane:** Administrators → CloudFront + ALB + WAF (Eval) or private ALB + WAF (Dual-AZ) → Admin UI
- **Storage plane:** EFS for customer share data; DynamoDB for ACL metadata; EFS paths for CTDB coordination artifacts

No customer file data traverses the Admin UI path.

4.7 Reference architecture and capacity

Capacity figures below are **planning references**, not measured SLOs. They are derived from FSLogix login characterization on real Windows Server clients, then LoginSim reproduction and CapacitySim per-node session characterization. These results are intended to support capacity planning for WorkSpaces and other FSLogix-based environments.

The 4-node Dual-AZ planning shape:

- Up to **20,000** concurrent users on **c6i.2xlarge**
- Up to **40,000** on **c6i.4xlarge** (always 2x)

- Linear scale to **8 nodes** (40,000 / 80,000)
- **2-node** requires the most headroom (a single failure doubles survivor load)
- **3 or more** nodes recommended for HA
- **9-16** Experimental (beta), not sized

Per-node login arrival rate (planning): **5 logins/s** on c6i.2xlarge and **10 logins/s** on c6i.4xlarge. A 30-minute login window doubles **completed** logins versus 15 minutes; concurrent seated users stay on the capacity table.




Figure 4-6 · 4-node Dual-AZ reference architecture

Supported concurrent users

Planning capacity from FSLogix characterization. c6i.4xlarge is 2x c6i.2xlarge.





Figure 4-7 · Supported concurrent users (planning capacity)

User logins per node

Arrival rate is constant. A 30-minute login window doubles completed logins, not concurrent seated users.



Cluster rate = nodes x per-node rate, linear through 8 supported nodes. Planning figures, not a LoginSim SLO.

8 x c6i.4xlarge: 80 logins/s - 72,000 (15 min) / 144,000 (30 min) completed logins

Figure 4-8 · User logins per node (planning figures)

5. Compute architecture

5.1 ECS cluster and capacity

RRSMB uses **Amazon ECS on EC2** with an Auto Scaling Group backing an ECS capacity provider. **1.3 supports 1-8 active cluster nodes**, configured at Marketplace launch (`ClusterSize`). **9-16** remains Experimental.

- **Eval Kit:** evaluation only (default 2 nodes)
- **2+ Dual-AZ nodes:** production high availability (4-node reference recommended; 3+ for maximum HA)

5.2 Replaceability and rolling updates

Each ECS task is **disposable**. RRSMB supports **automatic rolling updates** when newer Marketplace container images are published — scheduled for regional midnight by default on multi-node clusters.

5.3 Fencing and NLB health

Nodes that lose cluster authority are **fenced** from SMB service. NLB target group health checks steer traffic away from unhealthy or fenced tasks.

6. Storage architecture

6.1 Authoritative substrate: Amazon EFS

Customer file data lives on **customer-owned EFS** file systems and access points selected at share creation. RRSMB does **not** create customer share filesystems — you provision EFS in your account and attach shares via the Admin UI.

6.2 ACL and coordination state

NT ACL metadata is stored in **Amazon DynamoDB** in the customer account. CTDB uses EFS-backed coordination paths for recovery. This design keeps authoritative file bytes on EFS while preserving Windows ACL semantics over SMB.

6.3 Elastic Throughput and Infrequent Access

Amazon EFS Elastic Throughput follows login-storm and settle-tail demand instead of a 24x7 provisioned throughput tier sized for the Monday-morning peak. That is the durable cost distinction versus provisioned Windows file services. The size of the difference depends on peak-to-average ratio, retention, and the throughput tier the alternative would have to provision, so it is a property of each deployment rather than a fixed percentage.

Amazon EFS Lifecycle Management can move eligible files into Infrequent Access (and Archive) automatically while users keep the same SMB share namespace. Idle file data does not require a second file-server tier or a migration into another namespace. Both are Amazon EFS capabilities in the customer account.

7. Identity, security, and permissions

- **Active Directory is required** — domain join, Kerberos/NTLM, Windows-style ACLs
 - **Guest/anonymous SMB** is not supported
 - **Share permissions** in the Admin UI map to NT ACLs enforced on every SMB operation
 - **Admin UI** is HTTPS-only. Eval: CloudFront + ALB + WAF (internet-accessible). Dual-AZ Production: not internet-facing; private ALB + WAF only.
 - **SMB** access is restricted by `SmbClientCidr` on a private NLB
 - **Support Reports** offer redaction levels before sharing with vendor support
 - **Admission control** returns SMB2 `STATUS_SERVER_UNAVAILABLE` when a node is at capacity, instead of resetting the TCP connection
-

8. High availability and failure modes

Multi-node Dual-AZ deployments are designed to tolerate single-node and single-AZ failures through automatic CTDB failover and health-based NLB routing. Adding nodes (3–8) reduces exposure to compute and AZ faults and shortens the window in which the cluster runs without remaining redundancy. Effective availability depends on the customer's VPC, Active Directory, and

EFS design, and on the availability of underlying AWS services (Amazon EFS, DynamoDB, EC2). Roadrunner SMB does not publish a numeric SMB-service availability SLA at GA.

Validated failure modes include ECS task termination, EC2 host loss, AZ disruption (multi-node), and rolling replacement during software updates.

The Eval Kit is **not HA** and is intended for evaluation only. Dual-AZ is the only supported production configuration.

9. AWS Marketplace operations

9.1 Metered billing dimensions

Dimension	Meaning
ShareHours	Enabled, billable SMB shares per hour
ManagedStorageGBHours	Mounted EFS storage (GiB-hours) attached to billable shares

RRSMB software charges **stop** when the appliance is not running billable shares/storage. Delete the CloudFormation stack to stop Marketplace software metering.

Current rates are maintained on the [Roadrunner SMB website](#) — not in this architecture document.

9.2 What customers operate

Customers (or their MSP) manage:

- VPC, subnets, NAT, and endpoints (prerequisites)
- Active Directory connectivity
- EFS filesystems for shares
- Marketplace subscribe / unsubscribe
- Dual-AZ Admin reachability (VPN, Direct Connect, or jump box) and `SmbClientCidr` for SMB clients

RRSMB provides the Admin UI for shares, cluster health, billing status visibility, and support reports.

10. Architectural fit guidance

10.1 Ideal use cases

- VDI / FSLogix profiles and home directories on AWS
- Departmental team shares on elastic EFS
- Hybrid access over VPN or Direct Connect (SMB from allowed CIDRs)
- Cost-optimized SMB with EFS lifecycle tiering
- MSP-delivered file services in customer AWS accounts
- Because files live as standard files in Amazon EFS, Linux and Unix clients can access the same data over NFS while Windows users connect over SMB — one dataset, no second storage system. Ideal for build pipelines and data-processing estates alongside Windows shares. (Note: VDI/FSLogix profile workloads should remain SMB-only.)

10.2 When not to use RRSMB

- Database storage over SMB (SQL Server data/log volumes)
 - VSS-dependent backup workflows
 - DFS Replication as a native Windows target
 - Ultra-low-latency local block / HPC scratch
 - Workloads requiring unsupported SMB features (see Known Limitations in product documentation)
-

11. Roadmap (selected)

1.3 delivers login-scale hardening (8 NFS mux channels, accept-path admission, FSLogix login characterization on real Windows clients with LoginSim and CapacitySim) and a stronger Admin security posture (Eval CloudFront + ALB + WAF; Dual-AZ Production Admin not internet-facing). Supported cluster size is 1-8 (9-16 Experimental). Further work is prioritized from customer demand: Experimental 9-16 graduation, additional VDI ecosystems, and deeper admission telemetry.

12. Conclusion

Roadrunner SMB delivers **Elastic SMB for Amazon EFS — Windows-compatible SMB3 file shares** as an **AWS Marketplace appliance** — deployable in minutes, billed as **pay for hourly consumption use**, and operated entirely within the customer's AWS account.

By separating durable EFS storage from replaceable compute, enforcing consistent SMB and ACL behavior through clustered Samba and CTDB, and integrating natively with Active Directory, RRSMB provides a modern alternative to Windows file server sprawl on AWS — without proprietary hardware appliances or external control planes.

Get started: [AWS Marketplace](#) · [Roadrunner SMB — pricing & docs](#) · [Quick Start Guide](#)

Copyright © 2026 Roadrunner SMB, LLC. All rights reserved.

Roadrunner SMB and Elastic SMB are trademarks of Roadrunner SMB, LLC.

Patent pending.