

# Roadrunner SMB™

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

---

### Architecture & Design Whitepaper

Version 1.0 · June 2026

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

*Roadrunner SMB and Elastic SMB are trademarks of Roadrunner SMB, LLC. Patent pending.*

## 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 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. Most deployments reach a first working SMB share approximately 30 minutes after the stack completes — assuming an Active Directory domain is already accessible from the VPC.

RRSMB is particularly well suited for:

- VDI / FSLogix and user profile storage on AWS
- Departmental file shares and home directories
- MSP-managed AWS file services across multiple customer accounts
- Windows file server modernization with Active Directory and Windows ACL semantics preserved
- Multi-protocol workloads requiring both NFS and SMB access to the same dataset

### Availability and durability

RRSMB Dual-AZ deployments are designed for 99.99% SMB-service availability, consistent with the published AWS SLAs for every underlying component (EC2/ECS, NLB, EFS, and DynamoDB — each 99.99% for multi-AZ deployments). Customer file data is stored on Amazon EFS with 99.999999999% (11 nines) durability, inherent to EFS regional file systems that replicate across multiple Availability Zones. NT ACL metadata is stored in Amazon DynamoDB with multi-AZ replication and zero RPO. Backups via AWS Backup provide additional protection beyond these native guarantees.

Single-AZ / single-node deployments are for evaluation and non-HA use and do not carry the 99.99% availability claim.

For VDI-scale workloads, the economics of RRSMB versus Amazon FSx for Windows File Server (SSD Multi-AZ) are driven by two structural differences: EFS does not require pre-provisioned throughput capacity, and EFS Intelligent-Tiering / Infrequent Access can reduce storage costs for workloads with a meaningful proportion of idle files. The savings depend on your specific profile; see §6.3 for a worked scenario. Current Marketplace rates are at [roadrunnersmb.com](https://roadrunnersmb.com).

**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, Redis cache, or customer data egress is required.

**Not for every scenario:** RRSMB is not intended to replace every Windows file server workload. Workloads requiring VSS/shadow copies, DFS Replication, database storage over SMB, SMB Multichannel, or deep Windows kernel integration may be better served by 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 — FSLogix login storms generate hundreds of simultaneous opens, directory enumerations, and ACL evaluations within seconds. When underlying infrastructure is not optimized for these patterns, perceived performance degrades disproportionately. The traditional answer is over-provisioning: buy more throughput than steady-state requires, keep it running 24/7, pay for peak whether or not it materializes.

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.

Dimension	Roadrunner SMB	FSx for Windows File Server
<b>Underlying storage</b>	Amazon EFS (elastic, lifecycle tiering to IA/Archive)	Dedicated Windows file system (NTFS)
<b>SMB layer</b>	Clustered Samba on Amazon Linux 2023	Native Windows Server
<b>Data sovereignty</b>	All data and control plane in customer AWS account	AWS-managed service
<b>Multi-protocol (NFS+SMB)</b>	Yes — same EFS dataset over NFS and SMB simultaneously	SMB native; NFS requires separate architecture
<b>Throughput model</b>	Elastic — no pre-provisioning required; pay for actual I/O	Pre-provisioned throughput capacity (MB/s) billed whether used or not
<b>Cluster sizing</b>	1–4 nodes (1 for eval; 2+ for HA); roadmap: beyond 4	AWS-managed capacity model
<b>Active Directory</b>	Customer-managed AD (on-prem, EC2, or AWS Managed AD)	AWS Managed AD or self-managed AD
<b>VSS / shadow copies</b>	Not supported	Supported
<b>DFS Replication</b>	Not supported	Supported
<b>Availability (multi-AZ)</b>	99.99% (Dual-AZ, 2+ nodes)	AWS-managed; see FSx SLA
<b>File data durability</b>	99.999999999% (11 nines) via EFS regional replication	AWS-managed

<b>RRSMB software billing</b>	Hourly consumption: ShareHours + ManagedStorageGBHours; roadrunnersmb.com	FSx pricing (see AWS)
-------------------------------	--	-----------------------

**Choose Roadrunner SMB when:** you need EFS-backed elasticity without throughput pre-provisioning, the same EFS dataset over NFS and SMB, full in-account operational transparency, and pay-as-you-go hourly consumption on AWS Marketplace.

**Choose FSx for Windows File Server when:** you require native Windows features (VSS, DFSR) or 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. All components are standard AWS services. No proprietary appliance images, kernel modules, or external SaaS control planes are required or used.

### 3.2 Customer Data Sovereignty

All file data, ACL metadata, cluster state, and configuration remain in the customer's AWS account. RRSMB does not operate an external control plane, transmit customer file contents to any external system, or require persistent inbound network access to the customer's account. Customers retain full ownership and audit-ability of every resource RRSMB creates.

### 3.3 Replaceable Compute, Durable Storage

RRSMB nodes are replaceable ECS tasks. Customer file data resides on customer-owned EFS (11 nines durable). NT ACL metadata is stored in Amazon DynamoDB with zero RPO. Node replacement and rolling upgrades do not require storage migration, cluster quiesce, or administrator intervention.

### 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 — folder permissions, inherited ACLs, domain user/group assignments — without running Windows Server. NT ACL metadata is stored in DynamoDB via the `vfs_acl_ddb` VFS module, providing multi-AZ durability and zero-RPO recovery independent of node lifecycle.

### 3.5 Control Plane vs. Data Plane Separation

Administrative traffic (Admin UI, health, support reports) is logically and network-separated from SMB data traffic. A single Network Load Balancer handles both — TCP 445 for SMB and TLS 443 for Admin UI — with Security Group rules and separate CIDR parameters enforcing the boundary. No customer file data traverses the Admin UI path.

### 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. The acceleration layer operates in two tiers:

- **L1 Cache (node-local, RAM-backed):** Short-lived, best-effort semantic acceleration for frequently accessed metadata. Benchmark results demonstrate p50 cache hit latency below 5 microseconds and p90 below 10 microseconds, measured under 16-thread concurrent load on c6i.2xlarge against a representative FSLogix workload pattern (repeated open/stat/close on a directory tree of small files). Non-authoritative; TTL-based expiration with generation-counter invalidation ensures correctness.
- **L2 Authoritative Store (Amazon DynamoDB):** DynamoDB serves as the authoritative store for all Windows NT ACL metadata via the `vfs_acl_ddb` VFS module. Every ACL write commits to DynamoDB first. All nodes read from the same DynamoDB tables, enabling instant correct ACL access on any replacement node without restore steps. A hash-based deduplication design (two-table index + store) minimizes storage cost under inheritance-heavy ACL patterns common in large Windows file estates.

Small-file and metadata-heavy workloads — FSLogix profile containers, home directories, departmental shares — are the primary beneficiaries of the acceleration layer. Large sequential file transfers (backups, video, large datasets) are inherently throughput-bound and are served directly from EFS at EFS throughput rates.

### 3.7 Observability by Design

Integrated metrics (Prometheus sidecar), structured logging to CloudWatch, and Support Report generation from the Admin UI simplify day-2 operations and vendor support engagement.

## 4. Architectural Overview

RRSMB deploys into the customer's existing VPC via Marketplace CloudFormation. The stack performs VPC discovery and validation before creating product resources.

### 4.1 Buyer Flow

- Subscribe in AWS Marketplace
- Launch the CloudFormation template (VPC ID + node count) — approximately 10–15 minutes to `CREATE_COMPLETE`
- Complete First-Time Setup in the Admin UI (TLS certificate, domain join, first share) — approximately 10–15 minutes
- Mount shares from Windows clients on allowed networks

Most deployments reach a first working SMB share approximately 30 minutes after the stack completes. This assumes an Active Directory domain is already accessible from the VPC — RRSMB joins an existing domain; it does not create one. If AD connectivity is being established for the first time, allow additional time for DNS, VPN/Direct Connect, and domain controller reachability.

### 4.2 VPC Prerequisites

The CloudFormation stack validates the VPC before creating resources. Unsuitable VPCs are rejected immediately with a descriptive VPC Unsuitable: ... error. Prerequisites:

Requirement	Detail
Private subnets	At least one private subnet (Single-AZ) or two private subnets in different AZs (Dual-AZ). Nodes must not have public IPs.

<b>Public subnets</b>	At least one public subnet (Single-AZ) or two (Dual-AZ) for the Network Load Balancer. NLB handles both SMB and Admin UI.
<b>NAT Gateway</b>	Recommended for simplicity. ECS tasks need outbound connectivity for ECR image pulls, AWS API calls, and software update checks. Not required if all necessary VPC interface endpoints are configured (ECR API, ECR DKR, ECS, Secrets Manager, CloudWatch Logs) plus S3 and DynamoDB gateway endpoints. See <a href="https://roadrunnersmb.com/docs">roadrunnersmb.com/docs</a> for the full endpoint list.
<b>DNS resolution</b>	VPC DNS resolution and DNS hostnames must be enabled. DNS must resolve your Active Directory domain from within the VPC (Route 53 Resolver forwarding rules or DHCP options pointing to DC IPs).
<b>VPC endpoints</b>	S3 and DynamoDB gateway endpoints are recommended (no cost) and reduce NAT Gateway data transfer charges. When deploying without a NAT Gateway, additional interface endpoints are required — see VPC prerequisites documentation.
<b>Security Groups</b>	RRSMB creates its own security groups. Existing VPC security groups must permit the RRSMB SGs to reach Active Directory (see §7 for required AD ports).
<b>Active Directory reachability</b>	Domain controllers must be reachable from private subnets on Kerberos (TCP/UDP 88), LDAP (TCP 389), LDAPS (TCP 636), and DNS (UDP/TCP 53). See §7 for full list.

### 4.3 High-Level Block Diagram

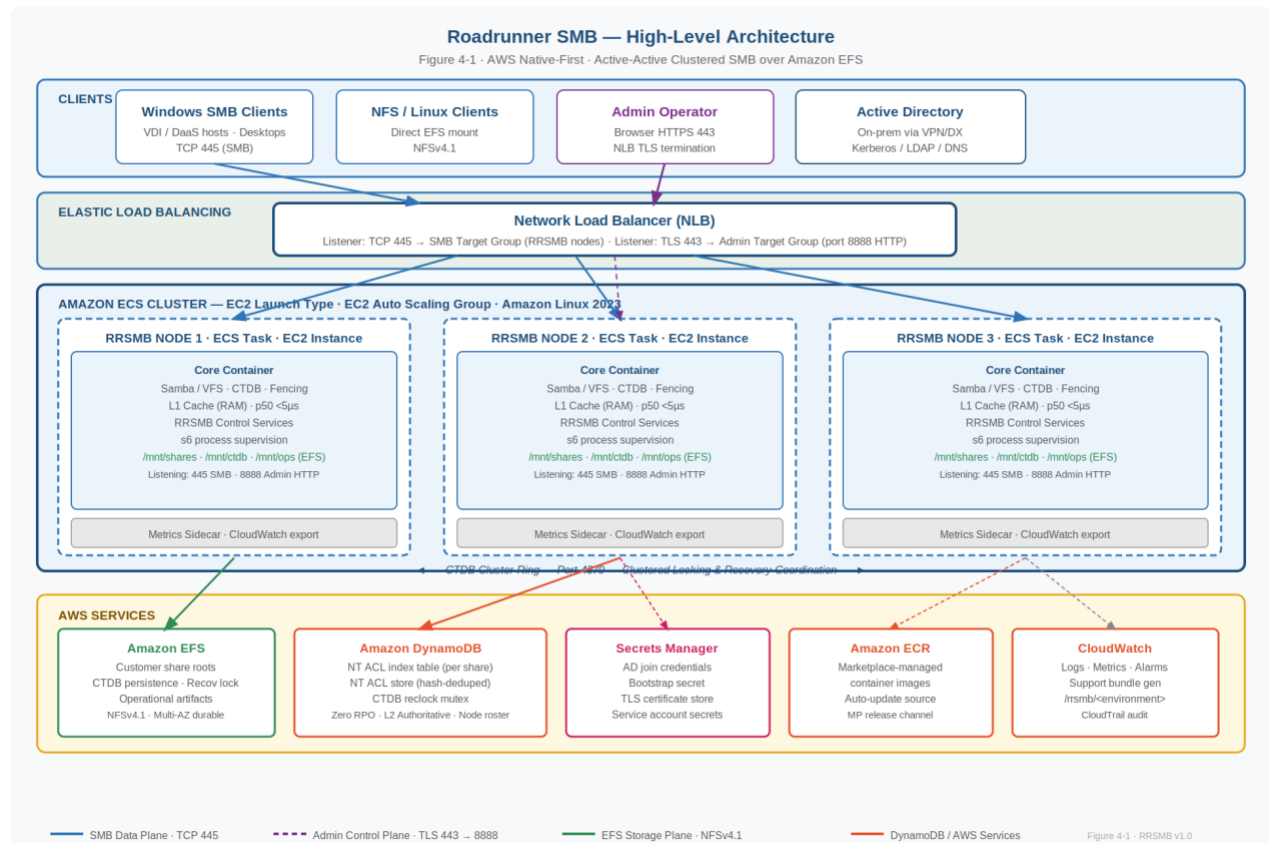


Figure 4-1 · High-level component view

### 4.4 Node Exploded View

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

- Samba / VFS — SMB3 protocol engine, NT ACL enforcement via vfs\_acl\_dbb (DynamoDB), Kerberos/NTLM authentication, winbind AD integration
- CTDB — clustered Samba state coordination, DynamoDB relock mutex, recovery orchestration
- Fencing — nftables host firewall, fail-closed on cluster authority loss, NLB health gate
- L1 Cache (RAM) — node-local metadata acceleration, p50 <5μs, TTL-based expiration
- RRSMB Control Services — domain join, Admin UI backend, cluster management API, share reconciler
- Metrics Sidecar — Prometheus metrics export to CloudWatch

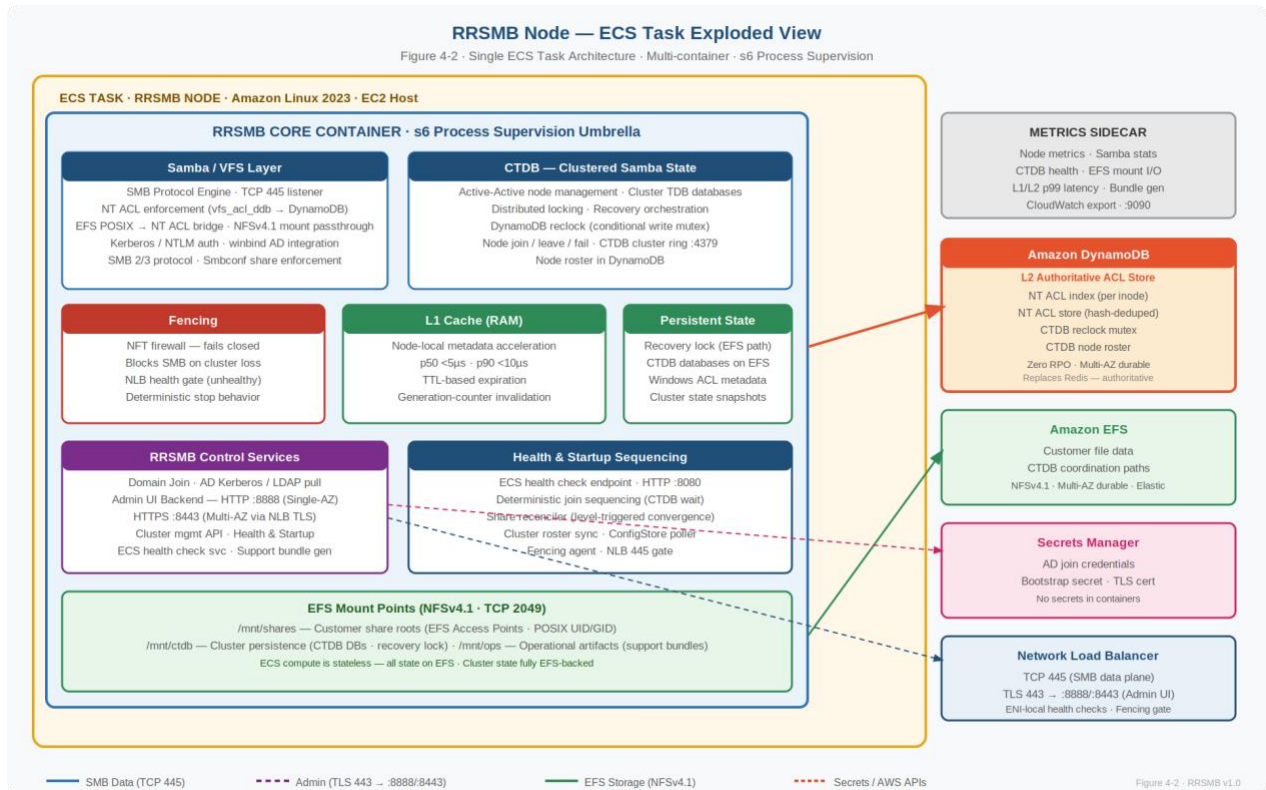


Figure 4-2 · RRSMB node — ECS task exploded view

### 4.5 Single-AZ vs Dual-AZ (HA)

	Single-AZ	Dual-AZ (HA)
Private subnets	1	2 (different AZs)
Public subnets	≥1 (NLB)	≥2
Default / typical nodes	1 (evaluation)	2+ (production HA)
AZ fault tolerance	None	Survives loss of one AZ

<b>Availability commitment</b>	Not HA; no SLA claim	99.99% (2+ nodes, both AZs)
<b>Use case</b>	POC, lab, dev/test	Production workloads

Marketplace delivery options: RoadrunnerSMB-SingleAZ and RoadrunnerSMB-DualAZ parent stacks with nested appliance stacks.

## 4.6 Networking Model

The Marketplace CloudFormation template uses a single Network Load Balancer (NLB) with two listeners:

Listener	Purpose	Access control
<b>TLS / 443</b>	Admin UI & First-Time Setup	AdminIngressCidr — defaults to open; restrict to VPN/Direct Connect CIDR for production
<b>TCP / 445</b>	SMB file sharing	SmbClientCidr — defaults to 10.0.0.0/8; set to your client network range

**Production hardening:** Restrict AdminIngressCidr to your VPN or Direct Connect address range. The Admin UI should never be internet-accessible in production — it exposes domain join credentials flow and share management. SmbClientCidr should be set to the specific subnets or CIDR ranges where Windows clients reside.

RRSMB nodes and EFS remain in private subnets. Outbound connectivity uses NAT Gateway (per AZ) for simplicity, or a full set of VPC interface and gateway endpoints for fully private deployments. S3 and DynamoDB gateway endpoints are recommended in either case (no cost, reduces data transfer charges).

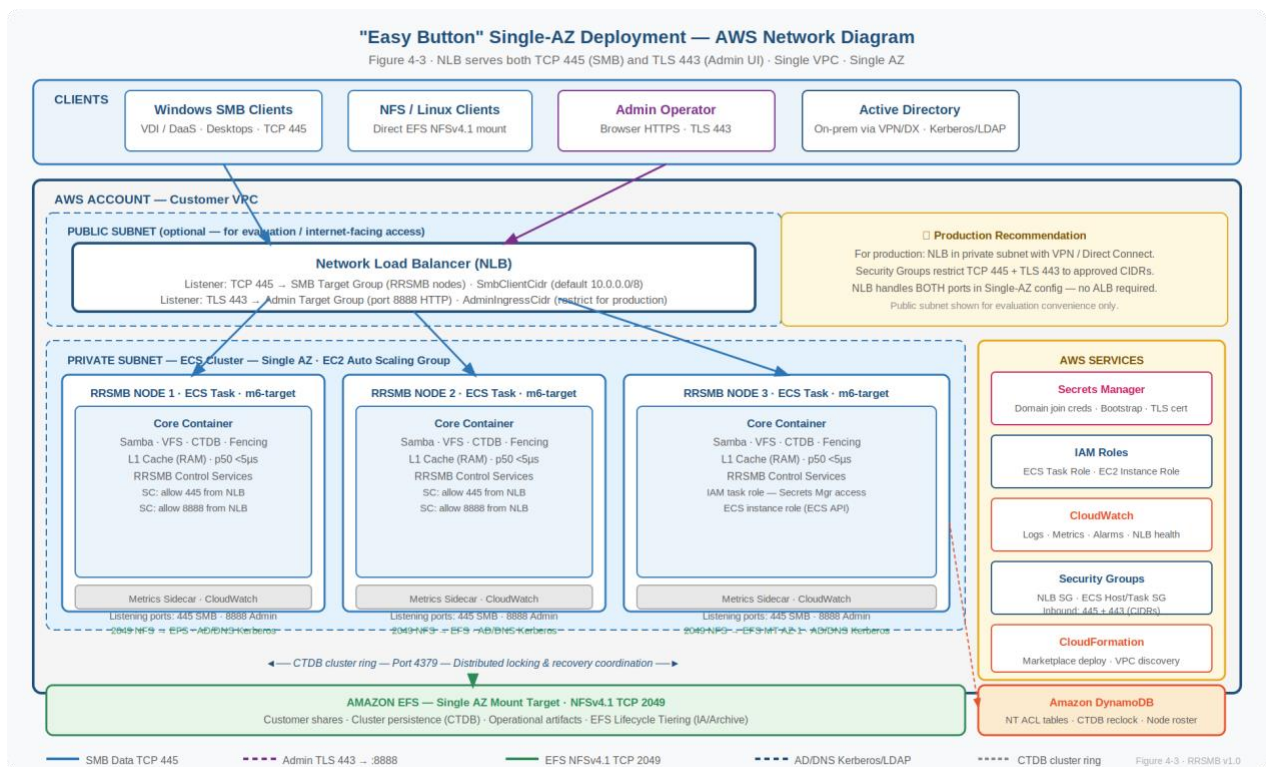


Figure 4-3 · Single-AZ deployment — AWS network diagram

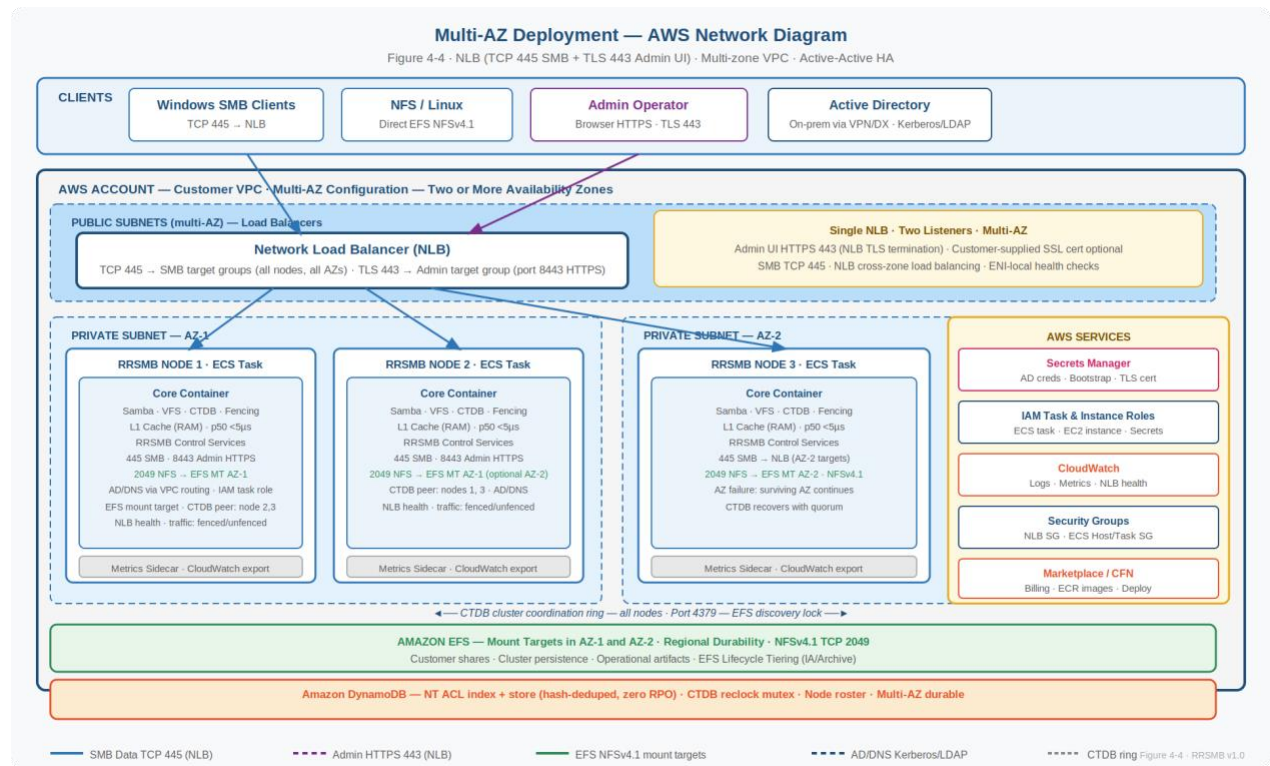


Figure 4-4 · Dual-AZ deployment — AWS network diagram

## 4.7 Control Plane vs. Data Plane

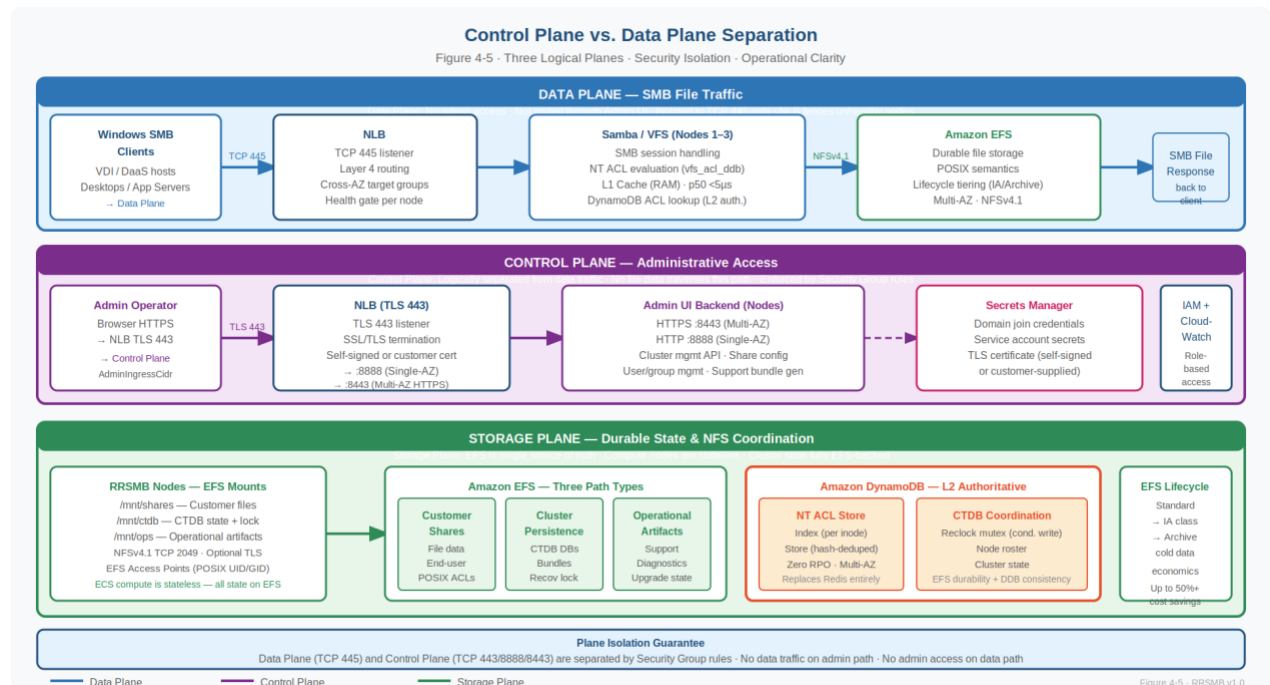


Figure 4-5 · Control plane vs. data plane separation

- Data plane: Windows SMB clients → NLB TCP/445 (SmbClientCidr) → Samba → EFS

- Control plane: Administrators → NLB TLS/443 (AdminIngressCidr) → Admin UI → Secrets Manager / IAM
- Storage plane: EFS for customer file data; DynamoDB for NT ACL metadata (L2 authoritative, zero RPO) and CTDB coordination; EFS paths for CTDB coordination artifacts

No customer file data traverses the Admin UI path.

## 5. Compute Architecture

### 5.1 ECS Cluster, Capacity, and Sizing

RRSMB uses Amazon ECS on EC2 with an Auto Scaling Group backing an ECS capacity provider. RRSMB is an appliance, not a microservices application — ECS provides exactly the right abstraction for this model: managed task placement and lifecycle, native integration with the NLB, ASG capacity, EFS mount targets, and IAM task roles, without the operational overhead of a Kubernetes control plane. RRSMB supports 1–4 active cluster nodes, configured at Marketplace launch via ClusterSize.

- 1 node: evaluation, lab, non-HA
- 2 nodes: production HA minimum — survives single node or AZ loss
- 3 nodes: recommended production — highest rolling-update availability; one node can be replaced while two serve
- 4 nodes: highest throughput / concurrent-user capacity

The 4-node limit is the current validated and supported maximum. It is not a hard architectural ceiling — the roadmap includes support for larger clusters based on customer demand. For workloads that will clearly outgrow 4 nodes, contact us before deploying.

Reference sizing (C6i/C6in family, EFS Elastic throughput — see §5.4 for instance selection guidance):

Tier	Concurrent Users	Nodes	EC2 Instance	AZ Config
<b>Evaluation</b>	Up to 50	1	c6i.large	Single-AZ
<b>Standard</b>	Up to 1,000	2–3	c6i.2xlarge	Dual-AZ
<b>High density</b>	Up to 5,000	3–4	c6in.2xlarge	Dual-AZ

#### Sizing note

These figures are reference estimates for FSLogix / home directory workloads with typical login-storm concurrency. Actual capacity depends on profile size, application I/O mix, and peak simultaneous login rate. Validate with DiskSPD or Login VSI against your specific workload before production cutover. RRSMB is network-intensive — every byte of SMB I/O also traverses the NIC as NFS I/O to EFS. c6i.2xlarge is the recommended instance for most deployments; c6in.2xlarge for high-density environments requiring higher per-node throughput.

### 5.2 Replaceability, Rolling Updates, and OS Patching

Each ECS task is disposable. RRSMB automatically detects newer approved Marketplace releases via a seller-published release manifest in Amazon S3. When a newer approved build is detected, a rolling update is scheduled for regional midnight by default — virtually eliminating user-visible SMB impact. Operators may apply updates immediately or defer via the Admin UI. Individual node replacement during a rolling update typically causes a brief (seconds) CTDB re-formation on the cluster; on 3+ node clusters this is transparent to SMB clients.

RRSMB nodes run on standard AWS ECS-optimized Amazon Linux 2023 AMIs — no custom or proprietary AMI is maintained. This is a meaningful operational and security advantage:

- OS-level security patches, kernel updates, and CVE fixes are delivered automatically. Node replacement during a rolling RRSMB update picks up the latest ECS-optimized AL2023 AMI, so the OS is refreshed along with the RRSMB application in the same operation.
- No separate OS patching workflow is required. Customers do not need to manage Linux kernel updates, security advisories, or AMI lifecycle for RRSMB nodes — it is handled as a natural consequence of the stateless, replaceable-task architecture.
- No custom AMI to maintain or audit. MSPs and enterprise customers with AMI compliance requirements benefit from RRSMB's use of AWS-published, AWS-maintained ECS-optimized images rather than a vendor-managed custom AMI with its own patching and validation cycle.

#### OS security posture

Rolling updates keep both the RRSMB application and the underlying Amazon Linux 2023 OS current simultaneously. The typical patch cycle is midnight-scheduled and transparent to SMB users on 2+ node clusters. For environments with strict patching SLOs, the Admin UI Check Now / Apply Now path allows operator-controlled immediate update application.

### 5.3 Fencing and NLB Health

Nodes that lose cluster authority are fenced from SMB service via nftables rules and NLB health-gate. Traffic is redistributed to healthy nodes within seconds. Windows SMB clients automatically reconnect through the NLB to a healthy node.

### 5.4 EC2 Instance Selection and Throughput Characteristics

RRSMB is a network-intensive workload, not a compute-intensive one. Every byte of SMB I/O traverses the EC2 network interface twice: once inbound as SMB traffic from the Windows client, and once outbound as NFS I/O to EFS. As a result, practical per-node throughput is bounded by approximately half the instance's total NIC bandwidth. CPU and memory are secondary considerations for most deployments.

The C6i and C6in families are recommended for RRSMB deployments because they provide high network bandwidth in a cost-efficient profile:

#### C6i family — up to 12.5 Gbps network bandwidth

Measured throughput on C6i instances: approximately 580 MB/sec read/write per node under a 1 MB sequential FIO benchmark. This reflects the double-traversal factor — 12.5 Gbps NIC capacity, with half available for effective file I/O after the SMB-in/NFS-out overhead is accounted for. The C6i family is the recommended standard choice for most RRSMB deployments.

Instance	vCPU	Network bandwidth	RRSMB estimated throughput/node
----------	------	-------------------	---------------------------------

<b>C6i.large</b>	2	Up to 12.5 Gbps	~580 MB/sec (measured, 1MB sequential FIO)
<b>C6i.xlarge</b>	4	Up to 12.5 Gbps	~580 MB/sec (same NIC ceiling as .large)
<b>C6i.2xlarge</b>	8	Up to 12.5 Gbps	~580 MB/sec (additional vCPU headroom for higher concurrency)

Note: all C6i sizes share the same 12.5 Gbps network ceiling. Larger C6i instances provide additional vCPU and memory for higher concurrent session counts, but do not increase per-node SMB throughput beyond the NIC limit. For higher throughput per node, use C6in.

### C6in family — high-throughput option

The C6in family is network-optimized and provides substantially higher NIC bandwidth. Throughput figures below are estimated by applying the same double-traversal model validated on C6i (measured throughput ≈ 46% of NIC bandwidth). These are estimates pending direct measurement; treat them as directional ceiling figures, not validated specifications.

Instance	vCPU	Network bandwidth	RRSMB estimated throughput/node
<b>C6in.large</b>	2	Up to 25 Gbps	~1.4 GB/sec (estimated)
<b>C6in.xlarge</b>	4	Up to 40 Gbps	~2.3 GB/sec (estimated)
<b>C6in.2xlarge</b>	8	Up to 50 Gbps	~2.9 GB/sec (estimated — recommended high-density instance)

#### C6in throughput — important caveats

These figures are estimated from the C6i measured result using the same NIC double-traversal model. They represent the NIC-bandwidth ceiling only. In practice, EFS per-client NFS throughput, Samba worker-thread concurrency, and CTDB coordination overhead will all impose constraints before the NIC ceiling is reached — particularly on the 2xlarge and larger sizes. Direct C6in measurement is in progress; this section will be updated with validated data.

c6in.4xlarge and larger instances are excluded from the table because EFS per-client throughput limits are expected to bind well before the NIC bandwidth at those sizes. Use c6in.2xlarge as the practical high-density ceiling until multi-mount-target or other EFS throughput expansion is validated.

Always verify current instance specifications at [aws.amazon.com/ec2/instance-types](https://aws.amazon.com/ec2/instance-types). Instance availability varies by region.

### Instance selection summary

c6i.2xlarge is the recommended instance for most RRSMB deployments — up to approximately 1,000 concurrent users per cluster (2–3 nodes). It provides the full 12.5 Gbps NIC bandwidth (measured ~580 MB/sec sustained throughput per node), 8 vCPUs for concurrent session handling, and proven performance under login-storm conditions.

c6in.2xlarge is recommended for high-density deployments — up to approximately 5,000 concurrent users per cluster (3–4 nodes). With an estimated ~2.9 GB/sec throughput per node from the higher-bandwidth NIC, c6in.2xlarge substantially expands per-cluster capacity. Validate EFS Elastic throughput headroom alongside NIC capacity at this scale — see the caveat above.

c6i.large is suitable for evaluation and single-node non-HA deployments where cost minimization outweighs throughput requirements.

## 6. Storage Architecture

### 6.1 Authoritative Substrate: Amazon EFS

Customer file data lives on customer-owned EFS file systems and access points. RRSMB does not create customer share filesystems — you provision EFS in your account and attach shares via the Admin UI. This means your data is never locked in a proprietary filesystem format, and you retain direct NFS access to the same data independently of RRSMB.

All nodes mount the same EFS filesystem concurrently via NFSv4.1. File data flows directly between Samba and EFS — there is no RRSMB-level file data cache. EFS Elastic throughput mode is recommended; Bursting mode can throttle under sustained high-concurrency metadata load.

#### Multi-protocol access (NFS + SMB)

Because files live as standard POSIX files on EFS, Linux and NFS clients can access the same data directly via NFSv4.1 mount while Windows clients access it via SMB — one dataset, no replication, no second storage system. This is particularly useful for build pipelines or data-processing estates that need to share data with Windows applications.

#### Multi-protocol consistency note

RRSMB enforces NT ACL semantics over SMB. Linux clients accessing the same files via NFS operate under POSIX permissions (UID/GID) which are independent of the NT ACLs. Files created via NFS will not have NT ACLs applied automatically — they receive synthesized NT ACLs based on POSIX permissions. For workloads requiring strict Windows ACL enforcement on all access paths, SMB-only access is recommended. VDI/FSLogix profile containers should always use SMB-only access.

#### Backup and data protection

EFS provides 11 nines file data durability natively through multi-AZ replication. For backup and point-in-time recovery, RRSMB share data is standard EFS and is fully compatible with AWS Backup. AWS Backup can create scheduled, policy-driven backups of EFS file systems with configurable retention. RRSMB does not include an integrated backup solution — customers should configure AWS Backup or their preferred backup tool against the EFS file system.

NT ACL metadata in DynamoDB is protected by DynamoDB's own multi-AZ replication with zero RPO. DynamoDB point-in-time recovery (PITR) can be enabled for additional ACL metadata protection.

### 6.2 NT ACL and Coordination State (DynamoDB)

Windows NT ACL metadata is stored in Amazon DynamoDB via the `vfs_acl_ddb` VFS module. DynamoDB is multi-AZ by design. A replacement node begins serving correct ACLs immediately on startup — no snapshot or restore steps required, zero RPO.

DynamoDB ACL storage uses a two-table deduplication design for cost efficiency at petabyte scale:

Table	Key Structure	Purpose
<b>rrsmb-acl-index</b> {cluster-id}	PK: ShareID + InodeID Attribute: aclHash	Lightweight per-inode ACL lookup index. Sub-millisecond reads; no ACL binary stored here.
<b>rrsmb-acl-store</b> {cluster-id}	PK: aclHash (SHA-256 of serialised SDDL) Attributes: aclBlob, refCount	Hash-deduplicated ACL binary storage. Identical inherited ACLs stored once. Enables petabyte-scale file estates at low ACL storage cost.

DynamoDB also serves as the CTDB distributed relock (conditional write mutex) and cluster node roster, providing AZ-resilient cluster coordination without EFS locking dependencies.

### 6.3 Lifecycle Economics — A Worked Scenario

The cost advantage of RRSMB versus FSx for Windows File Server depends on two structural factors: EFS does not require pre-provisioned throughput capacity, and EFS Infrequent Access storage class reduces cost for idle files. The following illustrative scenario shows the key cost drivers for a typical VDI deployment:

Illustrative scenario: 500-user VDI deployment, 10 TB total FSlogix profile storage, 5 shares

FSx for Windows File Server (SSD Multi-AZ):

Storage: 10 TB × \$0.23/GB/month = ~\$2,355/month

Throughput (pre-provisioned for login storm — 512 MB/s typical for 500-user VDI):  
~\$3,174/month

Approximate total FSx: ~\$5,529/month

Note: throughput capacity is paid whether or not it is used — login storm lasts minutes; you pay 24×7.

Roadrunner SMB on EFS (same 10 TB, 80% cold/IA | 20% active/standard, Elastic throughput):

EFS Standard (2 TB hot — active profiles): ~\$600/month

EFS IA (8 TB cold — idle profiles): ~\$128/month

EC2 (2 × c6i.2xlarge): ~\$490/month

NLB + NAT + DynamoDB: ~\$100/month

RRSMB software (5 shares, 10 TB): ~\$864/month

Approximate total RRSMB: ~\$2,182/month

Illustrative saving: ~61% (\$3,347/month) — from eliminating throughput pre-provisioning and EFS IA tiering.

5 shares reflects a realistic FSlogix deployment: FSlogix profiles, ODFC, and 2–3 departmental shares.

With more shares the software fee grows proportionally; fewer cold files reduces the IA saving.

Current RRSMB rates at [roadrunnersmb.com](http://roadrunnersmb.com). AWS service pricing from [aws.amazon.com](http://aws.amazon.com).

80% IA assumption reflects typical FSlogix environments where most profiles are inactive at any given time.

The throughput advantage is structural: FSx requires you to provision and pay for peak throughput capacity whether or not you use it. EFS Elastic throughput scales automatically and bills for actual I/O. For bursty workloads — login storms, morning peaks, month-end processing — this translates directly to cost.

## 7. Identity, Security, and Permissions

### 7.1 Active Directory Requirements

Active Directory is required for all RRSMB deployments. RRSMB joins an existing AD domain during First-Time Setup; it does not create or manage its own AD infrastructure.

Requirement	Detail
<b>AD functional level</b>	Windows Server 2008 R2 domain functional level or higher. Windows Server 2016/2019/2022 recommended.
<b>Supported AD types</b>	AWS Managed Microsoft AD (Simple or Standard), self-managed AD on EC2, on-premises AD reachable via VPN or Direct Connect.
<b>Domain join account</b>	A dedicated service account with permission to create computer objects in the target OU (or join rights on a pre-staged computer object). Domain Admin rights are not required.
<b>Admin UI account</b>	After First-Time Setup, any AD domain user can be granted Admin UI access. Domain Admins have Admin UI access by default (configurable in Settings).
<b>Kerberos / Kerberos time sync</b>	RRSMB nodes must be within 5 minutes of domain controller time (standard Kerberos requirement). AWS Time Sync (169.254.169.123) is used by default.
<b>DNS</b>	VPC DNS must resolve the AD domain FQDN. Configure Route 53 Resolver forwarding rules or set DC IP addresses in DHCP options. A DC IP override is available in the FTS wizard if DNS is not yet configured.

Required network ports from RRSMB node subnets to domain controllers:

Port	Protocol	Purpose
<b>53</b>	UDP/TCP	DNS — domain name resolution
<b>88</b>	UDP/TCP	Kerberos authentication
<b>389</b>	TCP	LDAP — directory queries, group membership
<b>636</b>	TCP	LDAPS — secure LDAP (optional but recommended)
<b>445</b>	TCP	SMB/RPC — domain join, AD replication checks

### 7.2 Authentication and Access Control

- Guest/anonymous SMB is not supported — all access requires valid AD credentials
- Share permissions in the Admin UI map to NT ACLs enforced on every SMB operation by Samba via `vfs_acl_ddb`

- Admin UI access is HTTPS-only (NLB TLS 443), restricted by AdminIngressCidr
- SMB access is restricted by SmbClientCidr
- A local break-glass account (rrsmb-admin) provides Admin UI access if AD is unreachable — set during First-Time Setup; store in a password vault

### 7.3 IAM and AWS Security Posture

RRSMB appliance runs under least-privilege IAM roles created in the customer's own account by the CloudFormation stack. No standing AWS credentials are supplied by the customer to RRSMB. The customer never hands RRSMB IAM admin rights — the stack creates its own scoped roles. The minimum IAM permissions required to deploy the stack are documented at [roadrunnersmb.com/docs/identity-and-permissions](https://roadrunnersmb.com/docs/identity-and-permissions).

### 7.4 Encryption

- Data in transit (SMB): Samba supports SMB encryption (SMB 3.x); enable per-share via `smb.conf` tuning if required by security policy
- Data in transit (NFS/EFS): EFS mount helper TLS option available for NFS-layer encryption
- Data at rest (EFS): EFS encryption at rest with AWS KMS (AWS-managed or customer-managed key) — configured at EFS filesystem creation, before attaching to RRSMB
- Data at rest (DynamoDB): DynamoDB encryption at rest with AWS-owned or KMS key
- Admin UI: TLS 443 via NLB — self-signed certificate provisioned during First-Time Setup; customer SSL certificate supported for production

### 7.5 Compliance and Licensing

Roadrunner SMB is built on Samba, which is open-source software licensed under the GNU General Public License v3 (GPLv3). Samba is widely used in enterprise environments including by large storage vendors. Customers do not receive source code obligations as end users. The RRSMB software layer built on top of Samba is proprietary. Contact us if your organization requires specific open-source licensing documentation for procurement review.

### 7.6 OS Patching and AMI Compliance

RRSMB nodes run on standard AWS-published ECS-optimized Amazon Linux 2023 AMIs. No custom or proprietary AMI is used. OS-level security patches and kernel updates are delivered automatically through the midnight rolling-update mechanism — node replacement refreshes both the RRSMB application and the underlying OS in a single operation. Environments with formal OS patching SLOs or AMI governance policies benefit from RRSMB's use of AWS-maintained images with no vendor-managed AMI lifecycle dependency.

## 8. High Availability and Failure Modes

Availability and durability commitments (Dual-AZ deployments)

99.99% SMB-service availability — consistent with the published 99.99% multi-AZ AWS SLAs for every underlying component: Amazon EC2/ECS, Elastic Load Balancing (NLB), Amazon EFS, and Amazon DynamoDB.

99.999999999% (11 nines) file data durability — inherent to EFS regional file systems, which replicate data across multiple Availability Zones automatically.

Zero RPO for NT ACL metadata — stored in DynamoDB with multi-AZ replication; a replacement node serves correct ACLs immediately without restore.

These figures are consistent with AWS published SLAs for the underlying services. RRSMB does not publish an independent SLA separate from the AWS service commitments on which it is built.

Multi-node Dual-AZ deployments are designed to tolerate single-node and single-AZ failures through automatic CTDB failover and health-based NLB routing. NLB health checks remove fenced or unhealthy nodes from rotation within seconds. Windows SMB clients reconnect automatically through the NLB.

Adding nodes (3–4) reduces exposure during rolling updates — two nodes can serve while a third is replaced. Effective availability depends on the customer's VPC, Active Directory, and EFS design, and on the underlying AWS service availability.

## 8.1 Validated Failure Modes

Failure mode	Behavior
<b>ECS task termination (single node)</b>	NLB removes unhealthy node; remaining nodes serve. ECS replaces the task; replacement converges full share state from EFS/DynamoDB within seconds on boot. No data loss.
<b>EC2 host failure</b>	ASG replaces the EC2 instance; ECS schedules a new task. SMB sessions reconnect via NLB. No data loss.
<b>AZ disruption (Dual-AZ, 2+ nodes)</b>	Nodes in surviving AZ continue serving. EFS and DynamoDB remain available (multi-AZ). CTDB re-forms with remaining nodes.
<b>Rolling software update</b>	Nodes replaced sequentially. 3+ node cluster: always two nodes serving. 2-node cluster: brief single-node window per update cycle. Midnight scheduling default minimizes client impact.
<b>CTDB cluster roster reformation</b>	Occasional (roughly 1-in-6 to 1-in-8 software update cycles) brief roster reformation during node replacement. Duration: typically under 60 seconds. Midnight scheduling makes this transparent in practice.
<b>AD unreachable (transient)</b>	Existing SMB sessions with cached Kerberos tickets continue. New connections requiring re-authentication may fail until AD is reachable. Admin UI remains accessible via local break-glass account.
<b>DynamoDB transient unavailability</b>	L1 cache serves recent ACL reads. Writes queue. ACL correctness maintained; brief latency increase possible. No data loss.

## 8.2 Disaster Recovery

RRSMB is a single-region deployment. There is no built-in cross-region failover at this time. For disaster recovery scenarios, the recommended approach is: (1) AWS Backup to replicate EFS data to a secondary region, and (2) redeploy the CloudFormation stack in the DR region pointing at the restored EFS. Because RRSMB's NT ACL metadata is in DynamoDB (not embedded in EFS), ACL data should be backed up separately via DynamoDB PITR and restored as part of any DR scenario. Cross-region replication and automated DR are roadmap items.

## 9. AWS Marketplace Operations

### 9.1 Metered Billing Dimensions

Dimension	Meaning
ShareHours	One billable unit per enabled SMB share per hour. A share is billable when it exists and is enabled — regardless of whether clients are actively connected.
ManagedStorageGBHours	The total capacity of the EFS file system mounted to RRSMB, in GiB-hours. This is the full EFS filesystem capacity, not the size of individual share folders.

**Important: billing is on total mounted EFS capacity, not individual share folder sizes**

If you attach a 50 TB EFS filesystem to RRSMB and expose a 5 TB share folder within it, billing is based on the full 50 TB EFS capacity. This mirrors how EFS itself bills — for the total storage provisioned across the filesystem, not by folder. Plan EFS sizing accordingly: for cost efficiency, use a dedicated EFS filesystem sized to your actual SMB workload rather than attaching a shared filesystem used for other purposes.

RRSMB software charges stop when the appliance is not running billable shares and storage. Delete the CloudFormation stack to stop Marketplace software metering. Your EFS data and DynamoDB ACL tables persist after stack deletion (EFS data until you delete the filesystem; DynamoDB tables are deleted with the stack by default). Current rates are at [roadrunnersmb.com](http://roadrunnersmb.com).

### 9.2 Software Updates

RRSMB automatically detects newer approved releases via a seller-published release manifest hosted in Amazon S3. The manifest is checked periodically and on-demand via Admin UI (Check Now). When a newer approved build is found, the update is scheduled for regional midnight by default; operators may apply immediately or defer via the Admin UI.

Software images are delivered through Marketplace-managed Amazon ECR and are subscription-gated. The release manifest gates which builds are offered for auto-update — only fully validated, Marketplace-approved builds are published. Buyer accounts do not need ECR DescribeImages permission; the manifest provides version information directly.

### 9.3 Upgrade Path from Earlier Builds

Existing running deployments automatically receive new approved releases through the midnight rolling-update mechanism. No manual stack update, re-launch, or migration is required for software updates. Share data, ACLs, and configuration persist across updates. For major version changes that require CloudFormation parameter updates, the release notes will document any required stack-level action.

### 9.4 What Customers Operate

- VPC, subnets, NAT, and endpoints (prerequisites, one-time setup)
- Active Directory connectivity and domain join account
- EFS filesystems for shares (created by customer; attached via Admin UI)
- AWS Backup configuration for EFS (recommended)
- Marketplace subscribe / unsubscribe

- Optional hardening of AdminIngressCidr and SmbClientCidr

RRSMB provides the Admin UI for share management, cluster health, billing status visibility, software update management, and support report generation. MSP deployments: each RRSMB installation runs in one customer VPC. There is no multi-tenant management plane; MSPs manage customer deployments individually or via their own automation against the RRSMB API.

## 10. Architectural Fit Guidance

### 10.1 Ideal Use Cases

- VDI / FSLogix profiles and home directories — the primary design target; metadata acceleration layer is optimized for this workload
- Departmental team shares on elastic EFS — pay for what you use, scale without re-provisioning
- Hybrid access over VPN or Direct Connect — SMB from approved client CIDR ranges
- Cost-optimized SMB with EFS lifecycle tiering for cold data
- MSP-delivered file services in customer AWS accounts — fully in-account, auditable, no shared control plane
- Multi-protocol workloads — Linux/NFS and Windows/SMB access to the same EFS dataset (see multi-protocol consistency note in §6.1)

### 10.2 When Not to Use RRSMB

- **Database storage over SMB:** SQL Server data/log volumes require low-latency block semantics that EFS/NFS cannot provide.
- **VSS-dependent backup workflows:** Volume Shadow Copy Service requires native Windows filesystem support.
- **DFS Replication targets:** DFSR requires native Windows replication infrastructure.
- **Ultra-low-latency local block / HPC scratch:** These require local NVMe or high-IOPS block storage.
- **Applications requiring SMB Multichannel:** Not supported. Single NLB connection per client session.

### 10.3 Common Questions

The following questions arise frequently in technical evaluations:

Question	Answer
<b>Does RRSMB work with on-premises Active Directory?</b>	Yes. RRSMB joins any AD domain accessible from the VPC — on-premises via VPN/Direct Connect, self-managed AD on EC2, or AWS Managed AD.
<b>What happens to my data if I uninstall RRSMB?</b>	Your EFS data persists independently. Deleting the RRSMB CloudFormation stack removes the appliance; it does not delete your EFS filesystem or its data. NT ACL metadata (DynamoDB) is removed with the stack by default. Back up ACL metadata via DynamoDB PITR before stack deletion if you plan to re-attach to a new RRSMB deployment.

<b>Can I use my existing EFS filesystem?</b>	Yes. RRSMB attaches to any EFS filesystem you specify. It does not format or restructure existing data.
<b>Is there a risk of vendor lock-in?</b>	Files are stored as standard POSIX files on EFS — accessible via NFS from any Linux client without RRSMB. NT ACL metadata lives in DynamoDB in your account. There is no proprietary filesystem format.
<b>Can multiple RRSMB deployments share one EFS?</b>	Not recommended. Each RRSMB cluster should use a dedicated EFS filesystem to avoid share-namespace conflicts and to ensure billing clarity.
<b>How do I size EFS throughput?</b>	Use EFS Elastic throughput mode (recommended). Elastic mode scales automatically with actual workload. Avoid Bursting mode for sustained high-concurrency login storms, which can exhaust burst credits.
<b>What is the maximum number of shares?</b>	No hard limit. Each share creates a DynamoDB ACL index. At very large share counts (thousands), contact us for guidance on DynamoDB provisioning.

## 11. Roadmap (Selected)

Roadmap priorities are driven by customer demand. The following areas are under active consideration:

- Cluster scaling beyond four nodes for high-user-density environments
- Login-storm pre-warming for VDI — proactive EFS metadata pre-fetch before peak login windows
- Cross-region disaster recovery automation
- Additional Admin UI hardening options including private-only Admin access modes
- Numeric availability SLA publication as operational data accumulates

Share feedback on roadmap priorities at [roadrunnersmb.com](https://roadrunnersmb.com) or through AWS Marketplace product feedback.

## 12. Conclusion

Roadrunner SMB delivers Elastic SMB for Amazon EFS — Windows-compatible SMB3 file shares as an AWS Marketplace appliance — deployable in approximately 30 minutes, billed as pay-as-you-go hourly consumption, and operated entirely within the customer's AWS account.

Dual-AZ deployments are designed for 99.99% SMB-service availability. Customer file data carries 11 nines EFS durability. NT ACL metadata in DynamoDB provides zero RPO. No Windows servers, proprietary appliance hardware, Redis caches, or external control planes are required.

By separating durable EFS storage from replaceable compute, enforcing consistent SMB and ACL behavior through clustered Samba and CTDB, and eliminating throughput pre-provisioning through EFS Elastic throughput, RRSMB provides a modern alternative to Windows file server sprawl on AWS — at a fraction of the operational complexity.

[Get started: AWS Marketplace](#) · [roadrunnersmb.com](#) · [Quick Start Guide](#)

Copyright © 2026 Roadrunner SMB, LLC. All rights reserved.  
*Roadrunner SMB and Elastic SMB are trademarks of Roadrunner SMB, LLC. Patent pending.*