IT Infrastructure Management Checklist for Growing SaaS Companies

Jared Knedler & Kyrylo Maznik February 4, 2026

Executive Summary

Infrastructure that worked at 10 customers breaks at 100, and rebuilds at 1,000 are expensive. This checklist describes the SRE-grade discipline a growth-stage SaaS needs to scale without rewrites: an SLO-driven observability stack, infrastructure as code, security baselines aligned to **CIS Benchmarks and NIST 800-53**, and the operational cadences that turn fragility into reliability. Customers running this discipline reach **99.95% measured uptime** within 60 days of onboarding and hold sev-1 MTTR under **15 minutes**.

Why SRE-Grade Infrastructure Matters

Enterprise customers write uptime into MSAs. SOC 2 reports document Availability as a Trust Services Criterion. Investors ask

about platform reliability in due diligence. The cost of a serious outage is no longer just refunded MRR; it is renewal risk and a hole in the next round's narrative.

The Google SRE handbook framed the modern operating model: *"Hope is not a strategy. Manage error budgets, not heroics."* The implication is operational: every customer-facing service has a Service Level Objective; every alert is tied to that SLO; every page is accountable to a measurable threshold. The discipline scales linearly; alert fatigue does not.

SLOs as the Operating Discipline

Define SLOs per customer-facing service, not per server. A typical pattern:

> **Availability SLO:** 99.9% over a 30-day rolling window for the customer-facing API; 99.95% for auth and billing.
>
> **Latency SLO:** p95 under 200ms for read endpoints, p95 under 500ms for write endpoints.
>
> **Error rate SLO:** below 0.1% 5xx responses over the same window.

The error budget — what the SLO permits — drives operational decisions. When the budget is healthy, ship faster. When the budget is depleted, slow down and stabilize.

Observability: The Datadog-Grafana-PagerDuty Stack

The exact tools matter less than the contract: traces, metrics, and logs in one queryable surface; alerts wired to error budgets; pages routed to on-call rotations tied to service ownership.

Instrumentation

Application instrumentation through **OpenTelemetry** avoids vendor lock-in at the SDK layer. Distributed traces flow into **Datadog APM** or **Grafana Tempo**; metrics into **Datadog Metrics** or **Prometheus**; logs into **Datadog Logs** or **Loki**. The combination matters less than the discipline of consistent instrumentation across services.

Burn-Rate Alerts

Burn-rate alerts page on error budget consumption velocity, not raw error counts. A 14.4x burn rate over 1 hour means the entire monthly budget burns in 2 hours; that is page-the-on-call territory. A 3x burn rate over 6 hours is a notification, not a page.

Dashboard Discipline

Every customer-facing service has a dashboard with: traffic, error rate, latency p50/p95/p99, saturation. The "four golden signals" from the Google SRE book remain the canonical starting set.

Incident Routing

PagerDuty rotations are tied to service ownership through tag-based routing. The on-call SRE acknowledges within 5 minutes; sev-1 incidents have a documented runbook linked from the alert; escalation chains are documented and tested.

Infrastructure as Code

The line between professional infrastructure operations and ad-hoc click-ops is whether the production environment can be reconstructed from version control.

Terraform as the Standard

Terraform with remote state in S3 plus DynamoDB locking is the operational baseline. Every AWS resource — VPC, subnet, security group, IAM role, RDS instance, S3 bucket, ALB, ASG, ECS service — is defined in code and reviewed through pull requests.

Module Organization

Reusable modules per resource family (network, data, compute, security) plus environment-specific compositions. The same module deploys to dev, staging, and production with environment-specific variable overrides — never separate copies of similar code.

State Hygiene

Production state isolated from non-production. State locking enforced. Drift detection runs nightly via **terraform plan** in CI; non-zero drift opens a ticket. Manual changes in the AWS console after week two of an engagement should be exception, not norm.

Account Hygiene

AWS Organizations with Service Control Policies that enforce regional restrictions and forbid certain destructive actions

AWS Control Tower for landing zone management at 5+ accounts

IAM Identity Center (formerly SSO) for human access; no IAM users for individuals

Tag policies that mandate the cost-allocation tag schema

Security Baselines

Security is encoded in the same Terraform modules that provision the resources, not added afterward.

CIS Benchmarks and NIST 800-53

The **CIS Benchmarks for AWS** provide concrete, hardenable controls per service. **NIST 800-53** defines the broader control catalog used as the basis for SOC 2 and FedRAMP control selection. Mapping CIS to NIST and to the AWS Well-Architected Framework's Security

Pillar gives one set of controls that satisfies multiple compliance regimes.

Three-Layer Hardening

Account guardrails: SCPs, IAM Identity Center, GuardDuty, Security Hub, AWS Config conformance packs

Workload hardening: CIS Benchmark images for EC2, SSM Patch Manager for OS-level patching, Inspector for vulnerability scanning

Application-level controls: WAF rules, ACM-managed TLS, Secrets Manager for runtime secrets, KMS for data-at-rest encryption

Continuous Verification

Controls evidenced in code are continuously verified by the same code. **AWS Config conformance packs** (or third-party GRC platforms) generate the audit evidence automatically. Quarterly threat-model review with engineering leads catches design-level gaps the rule engines miss.

Disaster Recovery and Resilience

Resilience patterns must match workload tier. Stateless customer-facing services run multi-AZ active-active with automatic failover; stateful services use cross-region read replicas and documented promotion runbooks.

RTO/RPO Targets

Tier 0 (customer-facing API + auth): RTO 15-60 minutes, RPO 0-5 minutes

Tier 1 (primary database): RTO 1-4 hours, RPO 1-15 minutes

Tier 2 (analytics, batch, internal tools): RTO 4-24 hours, RPO 1-24 hours

DR Drills That Run

Quarterly tabletops and annual full DR exercises produce dated evidence; per-component functional tests (read-replica promotion, ASG cross-region, S3 cross-region restore) run semi-annually. Without dated evidence, the DR program is a PDF.

Operational Cadences

24/7 on-call with documented sev-1 acknowledge time of 5 minutes

Weekly incident review (post-mortems for any sev-1 or sev-2)

Monthly cost-and-performance reports, runbook freshness audit

Quarterly architecture review, threat-model refresh, capacity planning, DR tabletop

Annually full DR exercise, security penetration test, compliance attestation cycle

Common Pitfalls

Per-server alerting without service SLOs: alerts on CPU, disk, and memory page on noise; alerts on customer-facing SLO breach page on actual customer harm.

Click-ops drift: manual AWS console changes accumulate until the Terraform plan is unmanageable. Lock down console write access for human users; require Terraform PRs.

Multi-AZ confused with multi-region: multi-AZ protects against AZ-level failure but not regional failure. Tier 0 needs multi-region.

Documented but untested DR: a runbook with no dated drill record is the most-cited SOC 2 finding. Run the drills, document the evidence.

Compliance theater: bolting on controls before audit week creates artifacts; encoding controls in Terraform creates an operational program.

Measured Outcomes

99.95% measured uptime within 60 days, often a 5-10x improvement over pre-engagement baseline

Sub-15-minute MTTR for sev-1 incidents

25-35% AWS cost reduction in the first quarter without performance regressions

SOC 2 + HIPAA evidence in code via Terraform modules and AWS Config — not screenshots before each audit

Zero hand-clicked changes in the AWS console after week two of engagement

How Jacobian Helps

Infrastructure rigor at growth-stage SaaS scale is a discipline most companies do not need to carry as headcount until they are well past Series B. We bring 10+ years of SRE experience across AWS, Azure, GCP, and hybrid environments, and we hand off a Terraform-managed environment with documented runbooks when you do hire your own team. The same Terraform modules that provision the infrastructure also generate the SOC 2, HIPAA, and ISO 27001 evidence your auditor needs — because our roots are in audit and compliance, not just engineering.

Resource Details

Author: Jared Knedler & Kyrylo Maznik

Published: February 4, 2026

Categories:

infrastructure SaaS IT management compliance

Download

Full Document

About This Resource

A practical checklist for managing IT infrastructure in growing SaaS companies. Covers compute, networking, security, automation, and compliance readiness.

Categories:

infrastructure SaaS IT management compliance