

SRE Maturity Assessment Guide

Instructions for Interactive and Offline Assessment

Bot Army Engineering | Assessment Toolkit v1.0

15

DOMAINS

75

QUESTIONS

450

MAX POINTS

30

MINUTES

5 MATURITY LEVELS

LEVEL	NAME	SCORE	AVG/DOMAIN
1	Ad-hoc	0-90	0-6
2	Foundational	91-180	7-12
3	Standardized	181-270	13-18
4	Advanced	271-360	19-24
5	Optimized	361-450	25-30

AFTER THE ASSESSMENT

- **Share:** Results with team and stakeholders
- **Identify:** Top 3 gaps (lowest scoring domains)
- **Review:** Rubrics for those domains
- **Read:** Improvement playbooks for guidance
- **Plan:** Create action items with owners

INTERACTIVE ASSESSMENT

- **Prepare:** Gather 2-4 team members, block 30-45 min
- **Start:** Open Interactive Assessment, enter team name
- **Answer:** Score each question 0-6 honestly
- **Review:** Check radar chart and domain breakdown
- **Export:** Save JSON for historical tracking

ASSESSMENT CADENCE

ACTIVITY	FREQUENCY
Full assessment	Quarterly
Progress review	Monthly
Action tracking	Weekly
Stakeholder report	Quarterly

OFFLINE PDF ASSESSMENT

- **Print:** Scoring worksheet + domain rubrics
- **Gather:** Same team prep as interactive
- **Score:** Use rubrics, record on worksheet
- **Calculate:** Sum domain totals (max 450)
- **Identify:** Circle domains below 13 points

COMMON MISTAKES

- Scoring aspirations instead of reality
- Ignoring evidence for high scores
- Rushing without team discussion
- Skipping domains that "don't apply"

SCORING GUIDE

SCORE	MEANING
0	Not practiced at all
2	Minimal/ad-hoc practice
3	Partial implementation
5	Strong implementation
6	Exemplary/industry-leading

BOT ARMY OWNERSHIP

TEAM	DOMAINS
SRE Bot	1, 6, 7, 8, 9, 11, 15
Ops Bot	4, 5, 12
Observability	2, 3
Security Bot	13
All Teams	10, 14

Measure to Improve

Target Level 3+ for all critical services.

Q1 SLOs & Error Budgets**Q1.1** How well-defined are your Service Level Indicators (SLIs)?

0 No SLIs defined | 2 Informal metrics tracked ad-hoc | 3 SLIs defined for some services | 5 Comprehensive SLIs for all critical services | 6 User-journey based SLIs with clear measurement methodology

Q1.2 How do you track and enforce error budgets?

0 No error budget concept | 2 Error budgets calculated but not enforced | 3 Error budgets tracked with manual reviews | 5 Automated burn rate alerts with policy enforcement | 6 Multi-window burn rates with automated feature freezes

Q1.3 How aligned are stakeholders on SLO targets?

0 No stakeholder awareness of SLOs | 2 Engineering aware, business not involved | 3 SLOs documented and shared with stakeholders | 5 Business and engineering co-own SLO targets | 6 SLOs embedded in contracts and product decisions

Q1.4 What happens when error budget is exhausted?

0 Nothing, we don't track error budgets | 2 Manual discussions, no formal process | 3 Documented escalation process | 5 Automatic feature freeze, reliability focus | 6 Proactive budget management prevents exhaustion

Q1.5 How do you review and iterate on SLOs?

0 SLOs never reviewed | 2 Reviewed when issues occur | 3 Quarterly reviews scheduled | 5 Regular reviews with customer feedback integration | 6 Continuous refinement based on user journey analysis

Q2 Observability**Q2.1** How comprehensive is your metrics coverage?

0 Ad-hoc metrics; no consistent approach across services | 2 System metrics (CPU/mem/disk) covered; app metrics ad-hoc | 3 RED/USE methods adopted for critical services | 5 Comprehensive coverage; consistent standards across services | 6 All services with golden signals; consistent naming/labels

Q2.2 How mature is your logging infrastructure?

0 Logs only on local disk, grep to debug | 2 Some centralized logging, unstructured | 3 Centralized structured logging with search | 5 Structured logs with correlation IDs and retention policies | 6 Intelligent log analysis with anomaly detection

Q2.3 How well do you implement distributed tracing?

0 No distributed tracing | 2 Tracing in some services, not correlated | 3 End-to-end tracing for critical paths | 5 Full tracing with service maps and latency analysis | 6 Continuous profiling integrated with tracing

Q2.4 How effective are your dashboards?

0 No dashboards or ad-hoc only | 2 Basic dashboards, often outdated | 3 Service-level dashboards maintained | 5 Golden signals dashboard per service with SLO tracking | 6 Self-service dashboard platform with templates

Q2.5 Can you correlate signals across metrics, logs, and traces?

0 Signals completely siloed | 2 Manual correlation via timestamps | 3 Some tooling for correlation | 5 Unified observability platform with correlation | 6 AI-assisted root cause analysis across signals

Q3 Alerting Strategy**Q3.1** What percentage of your alerts are actionable?

0 Unknown or mostly noise | 2 Less than 50% actionable | 350-80% actionable | 5 80-95% actionable, regular tuning | 6 >95% actionable, continuous improvement

Q3.2 How are alerts linked to runbooks?

0 No runbooks exist | 2 Some runbooks, not linked from alerts | 3 Runbooks linked for critical alerts | 5 All alerts link to runbooks, regularly updated | 6 Runbooks with automation hooks and versioning

Q3.3 How do you tune alert thresholds?

0 Set once, never tuned | 2 Tuned reactively after complaints | 3 Quarterly review of noisy alerts | 5 Data-driven tuning with noise metrics | 6 Automated threshold adjustment based on patterns

Q3.4 Do alerts correlate with SLO burn rates?

0 No SLO-based alerting | 2 Basic threshold alerts only | 3 Single-window burn rate alerts | 5 Multi-window burn rate alerts (fast + slow) | 6 Predictive alerting before budget exhaustion

Q3.5 How do you manage alert escalation?

0 No escalation process | 2 Informal escalation via chat/phone | 3 Documented escalation paths | 5 Automated escalation with on-call integration | 6 Intelligent routing based on context and expertise

Q4 Incident Response**Q4.1** How well-defined is your Incident Commander (IC) role?

0 No IC, whoever is available | 2 Informal IC, not always assigned | 3 IC role defined, rotation exists | 5 Trained ICs, clear handoff procedures | 6 IC certification program, regular drills

Q4.2 How do you track MTTD/MTTR?

0 Not tracked | 2 Ad-hoc calculations (spreadsheets, after-the-fact) | 3 Incident system with manual timestamp entry | 5 Automated capture from alerts and monitoring | 6 Real-time dashboards with trend analysis

Q4.3 How do you conduct postmortems?

0 No postmortems | 2 Ad-hoc reviews for major incidents | 3 Blameless postmortems with template | 5 Postmortems with tracked action items | 6 Learning reviews shared org-wide, patterns analyzed

Q4.4 How effective are your escalation paths?

0 No defined escalation | 2 Escalation exists but often unclear | 3 Documented escalation matrix | 5 Tested escalation with clear SLAs | 6 Automated escalation with fallback procedures

Q4.5 How do you train incident responders?

0 No training, learn by doing | 2 Informal shadowing | 3 Onboarding training exists | 5 Regular game days and tabletop exercises | 6 Certification program with continuous learning

Q5 On-Call Health**Q5.1** What percentage of time is spent on on-call work?

0 >50% of time on reactive work | 235-50% reactive | 325-35% reactive | 5 <25% reactive, rest proactive | 6 <15% reactive, highly automated

Q5.2 How many pages require response per on-call shift?

0 >10 pages per shift | 25-10 pages per shift | 32-5 pages per shift | 5 <2 pages per shift | 6 <1 page per shift; mostly proactive

Q5.3 How is on-call duty recognized and compensated?

0 No policy; on-call expected without recognition | 2 Informal; varies by manager or team | 3 Documented policy with time-off or pay | 5 Clear policy with pay, time-off, and flexibility | 6 Competitive compensation; on-call valued

Q5.4 How do you track on-call health metrics?

0 Not tracked | 2 Anecdotal feedback only | 3 Basic metrics (pages, hours) | 5 Comprehensive dashboard with trends | 6 Health metrics tied to improvement goals

Q5.5 How do you prevent burnout?

0 Burnout is common, no prevention | 2 React when people complain | 3 Rotation policies, some flexibility | 5 Proactive monitoring, load balancing | 6 Sustainable by design, team satisfaction high

Q6 Reliability Patterns**Q6.1** How do you implement circuit breakers?

0 No circuit breakers | 2 Ad-hoc implementation in some services | 3 Standard library used for critical paths | 5 All external calls protected, monitored | 6 Adaptive circuit breakers with auto-tuning

Q6.2 How standardized are timeouts and retries?

0 No timeouts or infinite waits | 2 Inconsistent timeouts across services | 3 Standard timeout policy documented | 5 Exponential backoff with jitter everywhere | 6 Context-aware adaptive timeouts

Q6.3 Do you isolate resources to prevent cascading failures?

0 No resource isolation; shared everything | 2 Some isolation, not systematic | 3 Critical services have dedicated resources | 5 Systematic isolation per dependency | 6 Dynamic isolation that adjusts to load

Q6.4 How do you handle graceful degradation?

0 All-or-nothing failures | 2 Some fallbacks, not systematic | 3 Degradation modes documented | 5 Automatic degradation with user communication | 6 Feature flags enable instant degradation

Q6.5 How do you prevent cascading failures?

0 Cascading failures happen regularly | 2 Some awareness, reactive fixes | 3 Load shedding for critical services | 5 Comprehensive protection at all layers | 6 Automatic blast radius containment

Q7 Capacity & Performance**Q7.1** How do you monitor utilization and saturation?

0 Not monitored | 2 Basic CPU/memory only | 3 USE method for critical resources | 5 Comprehensive USE dashboards with alerts | 6 Predictive capacity analysis

Q7.2 How mature is your autoscaling?

0 No autoscaling, manual only | 2 Basic CPU-based autoscaling | 3 Custom metrics-based autoscaling | 5 Predictive scaling with business signals | 6 ML-driven proactive scaling

Q7.3 How often do you load test?

0 Never or rarely | 2 Ad-hoc; only when issues arise | 3 Planned (scheduled or before releases) | 5 Automated in CI/CD | 6 Continuous with trend analysis

Q7.4 Do you have capacity models?

0 No capacity planning | 2 Gut feel; no documented approach | 3 Documented models for critical services | 5 Data-driven models updated regularly | 6 Automated capacity forecasting

Q7.5 How do you forecast demand?

0 No forecasting | 2 Ad-hoc estimates | 3 Historical trend analysis | 5 Integrated with business planning | 6 ML-based demand prediction

Q8 Release Engineering**Q8.1** How mature is your CI/CD pipeline?

0 Manual builds and deployments | 2 Basic CI, manual CD | 3 Full CI/CD for most services | 5 Standardized pipelines with quality gates | 6 Self-service platform with guardrails

Q8.2 How do you implement canary releases?

0 Big bang releases only | 2 Gradual rollout; ad-hoc monitoring | 3 Canary/baseline comparison; manual promotion | 5 Automated promotion based on SLOs | 6 Progressive delivery with auto-rollback

Q8.3 What are your DORA metrics?

0 Not tracked | 2 Low performer (monthly deploys, >60 lead time) | 3 Medium (weekly deploys, 1-6mo lead time) | 5 High (daily deploys, <1 week lead time) | 6 Elite (on-demand deploys, <1 day lead time)

Q8.4 How fast can you rollback?

0 Rollback not possible or hours | 2 230-360 minutes | 3 10-30 minutes | 5 <5 minutes, one-click | 6 Automatic rollback on SLO breach

Q8.5 How do you use feature flags?

0 No feature flags | 2 Ad-hoc flags in code | 3 Feature flag system for new features | 5 Comprehensive flag management with targeting | 6 Flags integrated with experiments and metrics

Q9 Tool & Automation**Q9.1** What percentage of time is spent on tool?

0 >50% tool | 235-50% tool | 325-35% tool | 5 <25% tool | 6 <10% tool, mostly engineering

Q9.2 How mature is your Infrastructure as Code?

0 ClickOps, manual provisioning | 2 Some IaC, not comprehensive | 3 IaC for new infrastructure | 5 Full IaC with GitOps workflow | 6 IaC with policy as code and drift detection

Q9.3 How much can developers do without waiting on ops?

0 Tickets required for everything | 2 Some self-service; most needs require tickets | 3 Common daily tasks are self-service | 5 Developer portal with golden paths | 6 Full platform with self-healing

Q9.4 How do you track and prioritize tool reduction?

0 Tool not tracked | 2 Anecdotal awareness | 3 Tool tracked, backlog exists | 5 Dedicated time for automation (20%) | 6 Tool elimination is a team OKR

Q9.5 How automated are routine operations?

0 Mostly manual | 2 Scripts exist, not maintained | 3 Key operations automated | 5 Comprehensive automation platform | 6 AI-assisted autonomous operations

Q10 Culture & Organization**Q10.1** How blameless are your postmortems?

0 Blame culture, people punished | 2 Lip service to blameless | 3 Genuinely blameless most times | 5 Blameless culture, focus on systems | 6 Failures celebrated as learning opportunities

Q10.2 Is psychological safety present?

0 Fear of speaking up | 2 Varies by team/manager | 3 Generally safe to raise concerns | 5 High safety, concerns welcomed | 6 Proactive seeking of diverse perspectives

Q10.3 How well do teams collaborate?

0 Siloed, competitive | 2 Collaboration when forced | 3 Regular cross-team interaction | 5 Embedded SREs, shared ownership | 6 Generative culture (Westrum)

Q10.4 How is knowledge shared?

0 Tribal knowledge, silos | 2 Documentation exists, outdated | 3 Regular knowledge sharing sessions | 5 Learning culture, communities of practice | 6 Organization-wide learning system

Q10.5 How is reliability ownership distributed?

0 Ops/SRE owns all reliability | 2 Developers aware but not responsible | 5 Shared on-call between dev and SRE | 5 You build it, you run it | 6 Reliability embedded in all teams

Q11 Chaos Engineering**Q11.1** How often do you run automated fault injection?

0 Never | 2 Ad-hoc; only after major incidents | 3 Planned experiments in staging | 5 Regular in staging; periodic in production | 6 Continuous automated chaos in production

Q11.2 How do you control blast radius?

0 No controls, hope for the best | 2 Manual safeguards | 3 Documented blast radius limits | 5 Automated abort conditions | 6 Intelligent blast radius with auto-scaling

Q11.3 Do you run game days or wargaming exercises?

0 No organizational exercises | 2 Informal tabletop discussions | 3 Annual scheduled exercises | 5 Quarterly with cross-team scenarios | 6 Regular full incident simulations

Q11.4 How do you apply learnings from chaos?

0 Findings ignored | 2 Ad-hoc follow-up | 3 Findings tracked, some fixed | 5 All findings tracked with SLAs | 6 Continuous improvement from chaos insights

Q11.5 What chaos tooling do you use?

0 No tooling | 2 Manual scripts | 3 Basic chaos tools (kill pods, etc.) | 5 Comprehensive platform (Gremlin, LitmusChaos) | 6 Custom platform integrated with observability

Q12 Disaster Recovery**Q12.1** Are your disaster recovery targets defined and validated?

0 No defined recovery targets | 2 Targets defined; never tested | 3 Targets defined; tested occasionally | 5 Regularly validated; meets targets | 6 Continuously validated; exceeds targets

Q12.2 How do you verify backups are restorable?

0 Never tested | 2 Only tested when issues occur | 3 Occasional restore tests | 5 Regular tests with data verification | 6 Automated continuous validation

Q12.3 How do you test failover?

0 Failover never tested | 2 Tested once, years ago | 3 Annual DR drills | 5 Quarterly failover tests | 6 Regular active-active failover

Q12.4 Do you have multi-region capability?

0 Single region only | 2 Cold standby in another region | 3 Warm standby with manual failover | 5 Hot standby with automated failover | 6 Active-active multi-region

Q12.5 How automated is recovery?

0 Fully manual, tribal knowledge | 2 Documented runbooks | 3 Partially automated | 5 Mostly automated with one-click recovery | 6 Self-healing with automatic recovery

Q13 Security Reliability**Q13.1** How do you manage secrets?

0 Secrets in source code or env vars | 2 Basic secrets storage, manual rotation | 3 Secrets vault with access control | 5 Dynamic secrets with auto-rotation | 6 Zero-trust secrets with audit logging

Q13.2 How are certificates managed?

0 Manual renewal, outages from expiry | 2 Calendar reminders for renewal | 3 Automated monitoring of expiry | 5 Auto-renewal (cert-manager, ACME) | 6 Short-lived certs with continuous rotation

Q13.3 How do you scan for vulnerabilities?

0 No scanning | 2 Ad-hoc scans | 3 Scheduled scans, manual remediation | 5 CI/CD integrated scanning with blocking | 6 Continuous scanning with auto-remediation

Q13.4 How often do you rotate credentials?

0 Never or when compromised | 2 Annually | 3 Quarterly | 5 Monthly or on-demand | 6 Continuous rotation (short-lived)

Q13.5 How do you handle security incidents?

0 No security incident process | 2 Ad-hoc response | 3 Security incident runbooks exist | 5 Dedicated security incident response team | 6 Automated detection and response (SOAR)

Q14 Documentation**Q14.1** How current is your architecture documentation?

0 No architecture docs | 2 Outdated diagrams | 3 Docs exist, updated occasionally | 5 Current docs, reviewed quarterly | 6 Auto-generated from code/infra

Q14.2 Do runbooks exist for all alerts?

0 No runbooks | 2 Runbooks for some alerts | 3 Runbooks for critical alerts | 5 All alerts have runbooks | 6 Executable runbooks with automation

Q14.3 How do you track architecture decisions?

0 No record of decisions | 2 Decisions in chat/email | 3 Some ADRs written | 5 ADR process followed consistently | 6 ADRs linked to code and searchable

Q14.4 How do you keep docs up-to-date?

0 Docs abandoned after creation | 2 Updated when someone notices issues | 3 Review cadence exists | 5 Docs as code in PRs | 6 Automated freshness checks

Q14.5 Can new team members onboard via docs?

0 Heavy reliance on shadowing | 2 Some docs, mostly tribal knowledge | 3 Onboarding guide exists | 5 Self-service onboarding possible | 6 Comprehensive onboarding with exercises

Q15 Dependency Management**Q15.1** Do you have a complete service map?

0 No service map | 2 Partial, outdated map | 3 Manual service map maintained | 5 Auto-discovered service map | 6 Real-time dependency graph with health

Q15.2 How do you track vendor SLAs?

0 Vendor SLAs unknown | 2 SLAs known but not monitored | 3 Major vendors monitored | 5 All vendors tracked with dashboards | 6 SLA compliance automated with alerts

Q15.3 How do you monitor dependency health?

0 No dependency monitoring | 2 Manual checks or vendor status pages | 3 Health endpoints monitored | 5 Comprehensive dependency dashboard | 6 Predictive dependency health analysis

Q15.4 What's your strategy for vendor outages?

0 No strategy, wait for vendor | 2 Manual workarounds | 3 Documented fallback procedures | 5 Automated failover to alternatives | 6 Multi-vendor redundancy by default

Q15.5 How do you manage library dependencies?

0 Dependencies not tracked | 2 Manual review occasionally | 3 Automated vulnerability scanning | 5 Automated updates with testing | 6 Dependency governance with policies

SRE Maturity Scoring Worksheet

Offline Assessment Form - Print and Complete

Bot Army Engineering | Assessment Toolkit

Team:

Date:

Assessors:

SCORING GRID (5 QUESTIONS PER DOMAIN, 0-6 POINTS EACH)

#	DOMAIN	Q1	Q2	Q3	Q4	Q5	TOTAL	LEVEL	NOTES
1	SLOs & Error Budgets						/30		
2	Observability						/30		
3	Alerting Strategy						/30		
4	Incident Response						/30		
5	On-Call Health						/30		
6	Reliability Patterns						/30		
7	Capacity & Performance						/30		
8	Release Engineering						/30		
9	Toil & Automation						/30		
10	Culture & Organization						/30		
11	Chaos Engineering						/30		
12	Disaster Recovery						/30		
13	Security Reliability						/30		
14	Documentation						/30		
15	Dependency Management						/30		

TOTAL SCORE: **/450** Maturity Level:

MATURITY LEVELS

LEVEL	NAME	SCORE
1	Ad-hoc	0-90
2	Foundational	91-180
3	Standardized	181-270
4	Advanced	271-360
5	Optimized	361-450

TOP 3 PRIORITY GAPS

- Domain: _____ Score: ___ Action: _____
- Domain: _____ Score: ___ Action: _____
- Domain: _____ Score: ___ Action: _____

QUESTION SCORING

SCORE	MEANING
0	Not practiced
2	Minimal/ad-hoc
3	Partial implementation
5	Strong implementation
6	Exemplary

NEXT STEPS

- Review rubrics for low-scoring domains
- Read improvement playbooks
- Create action items with owners
- Schedule quarterly reassessment

Domain 1: SLOs & Error Budgets

Service Level Objectives and Error Budget Management

SRE Bot / Foundations | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	No formal SLOs; availability discussed informally; no error budgets
2	Basic SLOs for some services; not consistently tracked; no budget enforcement
3	SLOs for critical services; error budgets calculated; basic burn rate monitoring
4	Comprehensive SLOs; budgets enforced; dev slowdowns when budget exhausted
5	SLOs drive all decisions; multi-window burn rates; automated freezes

ANTI-PATTERNS (RED FLAGS)

- Setting 100% availability targets (impossible, expensive)
- SLOs without error budget policies
- Engineering-only SLOs, no business alignment
- No consequence for budget violations
- Static SLOs that never evolve

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How well-defined are your SLIs?	6
2	How do you track/enforce error budgets?	6
3	How aligned are stakeholders on SLO targets?	6
4	What happens when error budget exhausted?	6
5	How do you review and iterate on SLOs?	6

EVIDENCE CHECKLIST

- SLO documentation exists and is up-to-date
- Error budget dashboards visible to stakeholders
- Historical SLO compliance data available
- Error budget policy with escalation process
- Evidence of SLO-driven prioritization decisions

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Observability	SLIs require metrics/logs infrastructure
Alerting	Burn rate alerts drive incident response
Release Eng	Error budgets gate feature releases

FOCUS AREAS

- **SLI Definition:** User-journey based indicators with clear measurement
- **SLO Targets:** Realistic, stakeholder-aligned availability goals
- **Error Budget Policy:** Clear consequences for budget violations
- **Stakeholder Alignment:** Business and engineering co-ownership

Error Budgets Enable Velocity

Managed risk, not zero risk.

Domain 2: Observability

Metrics, Logs, Traces, and Dashboards

Observability Bot | Observability | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL CRITERIA

- 1 Minimal logging; no centralized metrics; debugging via SSH
- 2 Basic metrics/logs; some dashboards; siloed per team
- 3 Centralized observability stack; standard dashboards; basic tracing
- 4 Full pillars (metrics, logs, traces); correlation; self-service
- 5 Exemplars, continuous profiling; AI-assisted analysis

ANTI-PATTERNS (RED FLAGS)

- Debugging production via SSH
- Metrics without context (no labels/tags)
- Logs without structured fields
- Dashboard sprawl with no ownership
- Observability as afterthought

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How comprehensive is your metrics coverage?	6
2	How mature is your logging infrastructure?	6
3	How well do you implement distributed tracing?	6
4	How effective are your dashboards?	6
5	Can you correlate across signals?	6

EVIDENCE CHECKLIST

- Centralized metrics platform (Prometheus, Datadog, etc.)
- Log aggregation with search capability
- Tracing enabled for critical paths
- Service-level dashboards exist
- Runbooks link to relevant dashboards

RELATED DOMAINS

DOMAIN	RELATIONSHIP
SLOs	SLIs derive from observability data
Alerting	Alerts query observability backend
Incidents	Dashboards critical for diagnosis

FOCUS AREAS

- **Metrics:** RED/USE methods, cardinality management
- **Logs:** Structured logging, centralized aggregation
- **Traces:** Distributed tracing, context propagation
- **Dashboards:** Service-oriented, actionable visualizations

Observe, Don't Guess

Data-driven debugging at scale.

Domain 3: Alerting Strategy

Actionable Alerts and Runbooks

Observability Bot | Observability | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Few alerts; mostly noisy; no runbooks; alert fatigue common
2	Basic alerts exist; high noise ratio; some documentation
3	SLO-based alerts; runbooks linked; regular tuning
4	Multi-window burn rates; <5% noise; automated tuning
5	Self-healing alerts; ML anomaly detection; proactive

ANTI-PATTERNS (RED FLAGS)

- Alerting on causes, not symptoms
- >20% non-actionable alerts
- No runbooks or outdated runbooks
- Alert storms during incidents
- Alerts ignored due to fatigue

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	What % of alerts are actionable?	6
2	How are alerts linked to runbooks?	6
3	How do you tune alert thresholds?	6
4	Do alerts correlate with SLO burn rates?	6
5	How do you manage alert escalation?	6

EVIDENCE CHECKLIST

- Alert actionability metrics tracked
- Runbooks exist for all critical alerts
- Alert noise ratio <20%
- Multi-window burn rate alerts configured
- Regular alert review cadence

RELATED DOMAINS

DOMAIN	RELATIONSHIP
SLOs	Burn rate alerts derive from SLOs
Observability	Alerts query observability data
On-Call	Alert quality affects on-call health

FOCUS AREAS

- **Actionability:** Every alert should have a clear action
- **SLO-Based:** Alert on error budget burn, not thresholds
- **Runbooks:** Documented response procedures
- **Tuning:** Regular noise reduction reviews

Alert on Symptoms

Every page should require human action.

Domain 4: Incident Response

Incident Command, Escalation, and Resolution

Ops Bot / Operations | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Chaotic response; no IC role; hero culture; no learning
2	Basic severity levels; some escalation paths; informal IC
3	Defined IC role; runbooks used; postmortems written
4	Trained ICs; MTTD/MTTR tracked; blameless culture
5	Incident learning system; automated mitigation; chaos drills

ANTI-PATTERNS (RED FLAGS)

- Hero culture (same person always responds)
- Blame-focused incident reviews
- No severity classification
- Postmortem actions never completed
- Escalation unclear or broken

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How well-defined is your IC role?	6
2	How do you track MTTD/MTTR?	6
3	How do you conduct postmortems?	6
4	How effective are escalation paths?	6
5	How do you train incident responders?	6

EVIDENCE CHECKLIST

- IC rotation schedule exists
- Severity levels defined with examples
- Escalation matrix documented
- Postmortem template in use
- MTTD/MTTR dashboards available

RELATED DOMAINS

DOMAIN	RELATIONSHIP
On-Call	On-call handles initial response
Alerting	Alerts trigger incident flow
Culture	Blameless culture enables learning

FOCUS AREAS

- **IC Role:** Clear ownership during incidents
- **Escalation:** Defined paths with contact info
- **Metrics:** MTTD, MTTR, incident frequency
- **Learning:** Blameless postmortems with actions

Incidents Are Learning Events

Every outage makes us stronger.

Domain 5: On-Call Health

Sustainable On-Call and Burnout Prevention

Ops Bot / Operations | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Burnout common; no rotation; >50% time reactive
2	Basic rotation; frequent paging; compensation unclear
3	Regular rotation; <25% time on-call; comp policy exists
4	<2 incidents/shift; health tracked; follow-the-sun
5	Proactive on-call; minimal paging; team satisfaction high

ANTI-PATTERNS (RED FLAGS)

- Same people always on-call
- >5 incidents per shift average
- No compensation for pages
- High turnover due to burnout
- On-call seen as punishment

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	What % of time is spent on-call work?	6
2	How many incidents per on-call shift?	6
3	Is on-call compensation clear?	6
4	How do you track on-call health metrics?	6
5	How do you prevent burnout?	6

EVIDENCE CHECKLIST

- On-call rotation schedule published
- Incidents per shift tracked (<2 target)
- Compensation policy documented
- Team satisfaction surveys conducted
- Handoff procedures documented

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Alerting	Alert quality affects paging load
Incidents	Incident volume drives on-call stress
Culture	Healthy culture supports on-call

FOCUS AREAS

- **Rotation:** Fair distribution, follow-the-sun if global
- **Workload:** <25% time, <2 incidents/shift target
- **Compensation:** Clear policy, time-off for pages
- **Health:** Burnout tracking, satisfaction surveys

Sustainable Operations

<25% time, <2 incidents/shift.

Domain 6: Reliability Patterns

Circuit Breakers, Retries, Timeouts, Bulkheads

SRE Bot | Resilience | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	No defensive patterns; cascading failures common
2	Basic timeouts in some services; retry logic ad-hoc
3	Circuit breakers for critical paths; standardized timeouts
4	Bulkheads, load shedding; graceful degradation
5	Adaptive patterns; self-healing; antifragile design

ANTI-PATTERNS (RED FLAGS)

- No timeouts (infinite waits)
- Retry storms (no backoff)
- All-or-nothing failures
- Cascading failures across services
- No graceful degradation paths

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How do you implement circuit breakers?	6
2	How standardized are timeouts/retries?	6
3	Do you use bulkhead isolation?	6
4	How do you handle graceful degradation?	6
5	How do you prevent cascading failures?	6

EVIDENCE CHECKLIST

- Circuit breaker library in use (Hystrix, resilience4j)
- Timeout policy documented
- Retry strategy with backoff implemented
- Load shedding mechanisms exist
- Graceful degradation tested

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Chaos Eng	Test patterns via chaos experiments
Dependencies	Patterns protect from dep failures
Capacity	Load shedding prevents overload

FOCUS AREAS

- **Circuit Breakers:** Fail fast when dependencies unhealthy
- **Timeouts:** Bounded wait times for all calls
- **Retries:** Exponential backoff with jitter
- **Bulkheads:** Isolate failure domains

Design for Failure

Assume everything will fail.

Domain 7: Capacity & Performance

USE Method, Load Testing, Autoscaling

SRE Bot | Resilience | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	No capacity planning; reactive scaling; no load testing
2	Basic monitoring; manual scaling; occasional load tests
3	USE method applied; autoscaling configured; regular tests
4	Capacity models; predictive scaling; continuous perf tests
5	ML-based forecasting; cost-optimized; real-time adaptation

ANTI-PATTERNS (RED FLAGS)

- Scaling only when pages fire
- No load testing before releases
- Unknown system limits
- Over-provisioned for "safety"
- No performance budgets

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How do you monitor utilization/saturation?	6
2	How mature is your autoscaling?	6
3	How often do you load test?	6
4	Do you have capacity models?	6
5	How do you forecast demand?	6

EVIDENCE CHECKLIST

- USE method dashboards for all services
- Autoscaling policies configured
- Load testing in CI/CD pipeline
- Performance regression tests exist
- Capacity planning documentation

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Observability	USE metrics from observability
Reliability	Load shedding at capacity limits
Release Eng	Perf tests gate releases

FOCUS AREAS

- **USE Method:** Utilization, Saturation, Errors
- **Load Testing:** Regular stress tests in CI/CD
- **Autoscaling:** Horizontal scaling with proper signals
- **Forecasting:** Demand prediction for planning

Know Your Limits

Measure, model, scale proactively.

Domain 8: Release Engineering

CI/CD, Canary Deployments, DORA Metrics

SRE Bot | Release | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Manual deployments; release days are stressful; no rollback
2	Basic CI; some CD; deployments weekly/monthly
3	Full CI/CD; canary deployments; DORA metrics tracked
4	Elite DORA metrics; automated rollback; feature flags
5	Continuous deployment; zero-downtime; progressive delivery

ANTI-PATTERNS (RED FLAGS)

- Manual deployments with scripts
- Big bang releases
- No rollback capability
- Releases require downtime
- DORA metrics unknown

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How mature is your CI/CD pipeline?	6
2	How do you implement canary releases?	6
3	What are your DORA metrics?	6
4	How fast can you rollback?	6
5	How do you use feature flags?	6

EVIDENCE CHECKLIST

- CI/CD pipeline fully automated
- Canary or blue-green deployments
- DORA metrics dashboard exists
- Rollback tested and documented
- Feature flag system in use

DORA ELITE TARGETS

METRIC	ELITE TARGET
Deploy Frequency	Multiple/day
Lead Time	<1 hour
MTTR	<1 hour
Change Fail Rate	<15%

FOCUS AREAS

- **DORA:** Frequency, lead time, MTTR, change fail rate
- **Canary:** Progressive rollout with auto-rollback
- **Feature Flags:** Decouple deploy from release
- **Rollback:** <5 minute recovery capability

Deploy Boring

Releases should be non-events.

Domain 9: Toil & Automation

Toil Reduction, Self-Service, Infrastructure as Code

SRE Bot | Release | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	>50% toil; manual everything; ticket-driven ops
2	Some automation; toil not measured; ad-hoc scripts
3	Toil <50%; IaC for infra; some self-service
4	Toil <30%; full IaC; developer self-service
5	Toil minimal; platform engineering; autonomous ops

ANTI-PATTERNS (RED FLAGS)

- Tickets for everything (ops as bottleneck)
- ClickOps in production
- Undocumented tribal knowledge
- SRE team is ticket queue
- No time allocated for automation

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	What % of time is spent on toil?	6
2	How mature is your IaC?	6
3	What self-service exists for developers?	6
4	How do you track/prioritize toil reduction?	6
5	How automated are routine operations?	6

EVIDENCE CHECKLIST

- Toil % tracked (<50% target)
- Infrastructure managed via IaC
- Self-service portal for common tasks
- Automation backlog exists
- Time explicitly allocated for automation

RELATED DOMAINS

DOMAIN	RELATIONSHIP
On-Call	Reduce pages via automation
Release Eng	CI/CD reduces deploy toil
Documentation	Automate runbook execution

FOCUS AREAS

- **Toil:** Manual, repetitive, automatable work
- **IaC:** Infrastructure defined in code (Terraform, Pulumi)
- **Self-Service:** Developer portals, golden paths
- **Automation:** Script → tool → platform progression

Automate the Boring

<50% toil, or push back.

Domain 10: Culture & Organization

Blameless Culture, Psychological Safety, Learning

All Teams / Culture / Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Blame culture; heroes celebrated; silos; fear of speaking up
2	Some awareness; lip service to blameless; inconsistent
3	Blameless postmortems; cross-team collaboration; learning
4	Generative culture (Westrum); psychological safety; innovation
5	Learning organization; failure celebrated; continuous growth

ANTI-PATTERNS (RED FLAGS)

- Looking for "who" not "what" failed
- Punishing people for incidents
- Information hoarding
- Fear of asking questions
- Reliability is "ops problem"

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How blameless are your postmortems?	6
2	Is psychological safety present?	6
3	How well do teams collaborate?	6
4	How is knowledge shared?	6
5	How is reliability ownership distributed?	6

EVIDENCE CHECKLIST

- Blameless postmortem template in use
- Psychological safety surveys conducted
- Cross-team collaboration examples
- Knowledge sharing sessions regular
- SRE embedded with dev teams

WESTRUM CULTURE TYPES

TYPE	CHARACTERISTICS
Pathological	Blame, silos, fear
Bureaucratic	Rules, turf, tolerance
Generative	Learning, sharing, inquiry

FOCUS AREAS

- **Blameless:** Focus on systems, not individuals
- **Safety:** Safe to report errors, ask questions
- **Westrum:** Generative vs pathological culture
- **Learning:** Continuous improvement mindset

Blame Systems, Not People

Psychological safety enables excellence.

Domain 11: Chaos Engineering

Game Days, Blast Radius Control, Failure Injection

SRE Bot | Resilience | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	No chaos practice; only learn from real outages
2	Occasional game days; manual failure injection
3	Regular chaos experiments; blast radius controlled
4	Continuous chaos in staging; production game days
5	Chaos in production daily; antifragile systems

ANTI-PATTERNS (RED FLAGS)

- Chaos without hypothesis
- No blast radius controls
- Chaos findings ignored
- Only chaos in staging
- Chaos as one-time event

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How often do you run chaos experiments?	6
2	How do you control blast radius?	6
3	Do you run game days?	6
4	How do you apply learnings from chaos?	6
5	What chaos tooling do you use?	6

EVIDENCE CHECKLIST

- Chaos experiment runbooks exist
- Game day schedule published
- Blast radius controls documented
- Chaos findings tracked and fixed
- Production chaos (with controls)

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Reliability	Validate patterns via chaos
DR	Test DR via chaos experiments
Incidents	Build muscle memory for response

FOCUS AREAS

- **Experiments:** Hypothesis-driven failure injection
- **Blast Radius:** Start small, expand gradually
- **Game Days:** Scheduled team resilience exercises
- **Tooling:** Chaos Monkey, Gremlin, Litmus

Break Things on Purpose

Find failures before they find you.

Domain 12: Disaster Recovery

RPO/RTO, Backup Testing, Failover Procedures

Ops Bot / Operations | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	No DR plan; backups untested; single region
2	Basic backups; DR plan exists but untested
3	RPO/RTO defined; backups tested; failover documented
4	Regular DR drills; automated failover; multi-region
5	Active-active; automated recovery; continuous DR testing

ANTI-PATTERNS (RED FLAGS)

- Untested backups
- Unknown RPO/RTO
- Single point of failure
- DR plan never tested
- Manual recovery procedures

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	Are RPO/RTO defined and met?	6
2	How often do you test backups?	6
3	How do you test failover?	6
4	Do you have multi-region capability?	6
5	How automated is recovery?	6

EVIDENCE CHECKLIST

- RPO/RTO documented per service
- Backup restoration tested quarterly
- Failover runbooks exist
- DR drills conducted annually
- Multi-region deployment (if applicable)

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Chaos	Chaos tests DR capabilities
Security	Backup encryption, access
Incidents	DR invoked during major incidents

FOCUS AREAS

- **RPO:** Recovery Point Objective (data loss)
- **RTO:** Recovery Time Objective (downtime)
- **Backups:** Regular testing, not just creation
- **Failover:** Tested, documented procedures

Plan for Failure

Test your backups, test your failover.

Domain 13: Security Reliability

Secrets, Certificates, Vulnerability Scanning

Security Bot / Governance / Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL

CRITERIA

- 1 Secrets in code; manual cert management; no scanning
- 2 Basic secrets vault; some cert automation; ad-hoc scans
- 3 Secrets rotated; cert auto-renewal; regular scanning
- 4 Zero-trust principles; scanning in CI; short-lived creds
- 5 Dynamic secrets; continuous compliance; automated remediation

ANTI-PATTERNS (RED FLAGS)

- Secrets in source control
- Long-lived credentials
- Manual certificate renewals
- No vulnerability scanning
- Security as afterthought

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How do you manage secrets?	6
2	How are certificates managed?	6
3	How do you scan for vulnerabilities?	6
4	How often do you rotate credentials?	6
5	How do you handle security incidents?	6

EVIDENCE CHECKLIST

- Secrets vault in use (HashiCorp, AWS SM)
- Certificates auto-renew (cert-manager)
- Vulnerability scanning in CI/CD
- Credential rotation policy documented
- Security incident runbook exists

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Release Eng	Security gates in CI/CD
DR	Secure backup storage
Documentation	Security runbooks needed

FOCUS AREAS

- **Secrets:** Vault, rotation, no hardcoding
- **Certs:** Auto-renewal, short expiry
- **Scanning:** SAST, DAST, dependency scanning
- **Zero Trust:** Verify explicitly, least privilege

Security as Reliability

Secure systems are reliable systems.

Domain 14: Documentation

Architecture Diagrams, Runbooks, ADRs

All Teams / Governance / Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Tribal knowledge; outdated docs; no runbooks
2	Some docs exist; quality varies; runbooks partial
3	Architecture documented; runbooks for critical paths
4	Docs as code; ADRs tracked; runbooks tested
5	Docs auto-generated; executable runbooks; always current

ANTI-PATTERNS (RED FLAGS)

- Knowledge only in people's heads
- Docs abandoned after creation
- Runbooks that don't work
- No architecture diagrams
- Decisions not recorded

EVIDENCE CHECKLIST

- Architecture diagrams exist and are current
- Runbooks linked from alert definitions
- ADR repository maintained
- Documentation review process exists
- Onboarding docs enable self-service

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	How current is your architecture documentation?	6
2	Do runbooks exist for all alerts?	6
3	How do you track architecture decisions?	6
4	How do you keep docs up-to-date?	6
5	Can new team members onboard via docs?	6

Bot Army Engineering | SRE Maturity Rubric [Prev: Security](#) | [Next: Dependencies](#) Domain 14 of 15

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Alerting	Alerts link to runbooks
Incidents	Runbooks aid response
Dependencies	Service maps document deps

FOCUS AREAS

- **Architecture:** C4 diagrams, service maps
- **Runbooks:** Linked from alerts, tested
- **ADRs:** Decision records with context
- **Freshness:** Regular review cadence

Docs as Code

If it's not documented, it doesn't exist.

Domain 15: Dependency Management

Service Maps, Vendor SLAs, Dependency Health

SRE Bot | Release | Max 30 Points

0-6

AD-HOC

7-12

FOUNDATIONAL

13-18

STANDARDIZED

19-24

ADVANCED

25-30

OPTIMIZED

SCORING CRITERIA BY LEVEL

LEVEL	CRITERIA
1	Unknown dependencies; surprise failures from vendors
2	Partial dependency list; some vendor tracking
3	Service maps exist; vendor SLAs tracked; alerts on deps
4	Dependency health dashboard; degradation strategies
5	Auto-discovery; vendor SLA enforcement; antifragile design

ANTI-PATTERNS (RED FLAGS)

- Unknown critical dependencies
- No vendor status monitoring
- Single vendor for critical path
- Outdated library dependencies
- No fallback for vendor outage

ASSESSMENT QUESTIONS

#	QUESTION	MAX
1	Do you have a complete service map?	6
2	How do you track vendor SLAs?	6
3	How do you monitor dependency health?	6
4	What's your strategy for vendor outages?	6
5	How do you manage library dependencies?	6

EVIDENCE CHECKLIST

- Service dependency map exists
- Vendor SLAs documented and monitored
- Dependency health dashboard available
- Degradation strategies for critical deps
- Library dependency scanning automated

FOCUS AREAS

- **Service Maps:** Visual dependency graphs
- **Vendor SLAs:** Tracked, compared to internal SLOs
- **Health:** Dependency health as metric
- **Degradation:** Graceful handling of dep failures

RELATED DOMAINS

DOMAIN	RELATIONSHIP
Reliability	Circuit breakers for deps
Observability	Track dependency metrics
Security	Library vulnerability scanning

Know Your Dependencies

Your SLO is bounded by theirs.