ID	Domain	Control Name	Control Requirement	How to comply	Evidence (optional)	Evidence type	Source	Vendor Attestation (1/0)	Evidence link/location	Auditor notes
CS-1	Clinical Safety & Effectiveness	Validated clinical performance	Your model meets or exceeds an appropriate clinical benchmark for its intended use and population.	1) Define intended use and target population. 2) Prepare ground-truth test set independent of training. 3) Calculate primary metrics (e.g., sensitivity, specificity, AUC, PPV/NPV) and compare to current standard of care. 4) Document results and sign-off by a clinical lead.	Validation report including: intended use; dataset description (size, timeframe, sites); ground-truth method; inclusion/exclusio	Validation report (PDF)	CHAI, FDA			
CS-2	Clinical Safety & Effectiveness	Safety risks identified & mitigated	You have mapped patient safety failure modes and implemented mitigations.	1) Run a failure modes & effects analysis (FMEA) or similar. 2) Define mitigations (e.g., confidence thresholds, human review for edge cases, fail-safe defaults). 3) Assign owners and due dates. 4) Re-test high-risk scenarios.	Risk register listing top risks, severity, mitigations, owners, dates; proof of implemented mitigations (screenshots/configs); re-t	Risk register + screenshots	CHAI, NIST			
CS-3	Clinical Safety & Effectiveness	Post-deployment safety monitoring	You track and respond to AI-related incidents in real use.	1) Publish an incident reporting path for users. 2) Log incidents/near-misses and triage within defined SLAs. 3) Perform root-cause analysis and corrective actions. 4) Share safety advisories with customers as needed.	Safety monitoring SOP; sample redacted incident log; example RCA & corrective action; communication template.	SOP + logs	FDA, CHAI			
DI-1	Data Integrity & Lineage	Data provenance documented	You can trace where training/validation data came from and how it was processed.	1) Inventory all datasets with source system, date range, sites. 2) Record preprocessing/labeling steps and tools. 3) Note synthetic/augmented data use. 4) Store records under version control.	Data provenance register with dataset IDs, sources, preprocessing steps, labelers, and version references.	Provenance register (CSV/PDF)	NIST			
DI-2	Data Integrity & Lineage	Data quality controls	You check input/training data for completeness and correctness.	1) Define quality checks (completeness, range, outliers). 2) Run checks and log issues. 3) Remediate or exclude bad data. 4) Re-run checks and archive reports.	Most recent data quality report showing checks performed, issues found, remediation, and sign-off.	Quality report (PDF)	NIST			
FA-1	Fairness & Non-Discrimination	Subgroup performance tested	You evaluate model performance across key demographic groups.	1) Define subgroups (e.g., race, sex, age, language) relevant to use. 2) Compute metrics by subgroup. 3) Flag gaps exceeding pre-set thresholds. 4) Document results and actions.	Fairness test report with subgroup metrics, thresholds, gaps, and remediation plan if needed.	Fairness report (PDF)	CHAI, FDA			
FA-2	Fairness & Non-Discrimination	Bias mitigation applied	You took steps to reduce identified bias in data/model.	1) Review dataset representation; identify proxies. 2) Apply mitigation (re-weighting, re-sampling, debiasing models). 3) Re-test fairness. 4) Document changes and outcomes.	Bias mitigation memo describing methods used, before/after metrics, and approval.	Mitigation memo (PDF)	CHAI, NIST			
FA-3	Fairness & Non-Discrimination	Ongoing fairness monitoring	You have a plan to check fairness in production.	1) Schedule periodic subgroup analyses. 2) Track drift in subgroup distributions. 3) Escalate if gaps exceed thresholds. 4) Update model or workflow as needed.	Monitoring plan with cadence, metrics, thresholds, roles; sample dashboard or report.	Monitoring plan (PDF)	CHAI			
GO-1	Governance & Oversight	Defined AI governance	Roles and decision rights are formalized across the AI lifecycle.	1) Appoint accountable owners (clinical, technical, privacy, security). 2) Establish review checkpoints. 3) Publish RACI for development, deployment, monitoring. 4) Approve in governance charter.	Governance charter with roles, RACI, and stage-gate process; approval record.	Charter (PDF)	NIST, CHAI			
GO-2	Governance & Oversight	Change management	Model/data changes follow a documented review and approval process.	1) Submit change requests with impact analysis. 2) Perform risk/performance regression checks. 3) Obtain approval before release. 4) Maintain change log and rollback plan.	Change control SOP and a sample approved change request with test results.	SOP + sample CR	NIST			
GO-3	Governance & Oversight	Incident management & accountability	There is a clear path to pause or stop unsafe AI.	1) Define incident severity levels and SLAs. 2) Assign on-call roles and escalation. 3) Practice drills and document learnings. 4) Authority documented to suspend AI use.	Incident response plan with roles/SLAs; drill summary; suspension authority statement.	IR plan (PDF)	NIST, CHAI			
LR-1	Legal & Regulatory Compliance	Privacy & data protection compliance	You comply with applicable privacy laws (e.g., HIPAA, GDPR) for all data uses.	1) Map data flows and PHI handling. 2) Apply minimum necessary, de-ID where possible. 3) Ensure BAAs/DPAs signed. 4) Complete privacy review and approval.	Privacy compliance checklist; signed BAA/DPA samples; data flow diagram.	Compliance pack (PDF)	HIPAA/GDPR, CHAI			
LR-2	Legal & Regulatory Compliance	Regulatory status determined	You have confirmed if the AI is a regulated device and acted accordingly.	1) Classify intended use vs. regulatory definitions. 2) If regulated, provide clearance/approval. 3) Use only within cleared indications. 4) Maintain regulatory file.	Regulatory assessment memo; clearance/cert letter if applicable.	Regulatory memo (PDF)	FDA			
LR-3	Legal & Regulatory Compliance	Consent & disclosure	Patients/providers are informed when AI is used and data is processed.	1) Provide patient/provider notices. 2) Obtain consent where required. 3) Publish limitations and intended use. 4) Keep records of disclosures.	Standard disclosure text; consent form template; record of use sample.	Disclosure/consent (PDF)	CHAI, NIST			
VM-1	Model Validation & Monitoring	Pre-deployment validation completed	The model passed predefined acceptance criteria before go-live.	1) Define acceptance thresholds. 2) Validate on holdout/external data. 3) Document results vs thresholds. 4) Obtain approval to deploy.	Validation protocol & report showing thresholds and pass/fail with approvals.	Validation pack (PDF)	FDA, CHAI			
VM-2	Model Validation & Monitoring	Performance guardrails set	Minimum performance is defined and enforced in production.	1) Set metric floors (e.g., min AUC, sensitivity). 2) Configure alerts when below floor. 3) Define auto-pause or review process. 4) Document responsibilities.	Performance specification; ops playbook showing triggers and actions.	Spec + playbook	NIST, CHAI			
VM-3	Model Validation & Monitoring	Production monitoring & drift detection	You track real-world performance and data drift.	1) Implement dashboards for key metrics. 2) Monitor data distributions. 3) Investigate significant shifts. 4) Retrain/recalibrate as needed.	Screenshot or export of monitoring dashboard; drift report example.	Dashboard + report	NIST, CHAI			
SR-1	System Reliability & Resilience	High availability & recovery	Service reliability and recovery are planned and measured.	1) Define SLA uptime targets. 2) Implement backups and restoration tests. 3) Document RTO/RPO. 4) Monitor uptime and report breaches.	SLA document; DR plan; last restoration test report; uptime report.	SLA/DR pack	NIST			
SR-2	System Reliability & Resilience	Fail-safe behavior	If the AI is unavailable or uncertain, it fails safely.	1) Define safe defaults (e.g., human-only workflow). 2) Implement confidence thresholds. 3) Test failover paths. 4) Document behavior in user guide.	Design doc/screenshots showing fail-safe; test results of failover.	Design + test evidence	NIST, CHAI			
SR-3	System Reliability & Resilience	Stress & edge-case testing	You have tested robustness under load and unusual inputs.	1) Define load and edge scenarios. 2) Run stress tests. 3) Record errors/latency and fixes. 4) Re-test after fixes.	Stress test plan and results with remediation notes.	Test report (PDF)	NIST			
IM-1	Impact Measurement & ROI	Impact on outcomes or operations	You can show meaningful improvement attributable to the AI.	1) Define KPIs pre/post. 2) Collect baseline and post-implementation data. 3) Attribute impact (A/B or time-series). 4) Document results and limitations.	Impact report with KPI definitions, methods, results, and attribution.	Impact report (PDF)	CHAI			
IM-2	Impact Measurement & ROI	Adoption & feedback tracked	You monitor whether users adopt the tool and what they think.	1) Track usage/adoption metrics. 2) Collect structured feedback. 3) Act on feedback with changes. 4) Close the loop to users.	Adoption dashboard export; feedback summary and action log.	Dashboard + summary	CHAI			
IM-3	Impact Measurement & ROI	No undue workflow burden	Net workflow burden is acceptable and documented.	1) Identify potential burdens (alerts, clicks). 2) Measure time/steps added. 3) Mitigate or justify with benefits. 4) Get clinical sign-off.	Workflow impact assessment with measures and mitigation; sign-off.	Assessment (PDF)	CHAI			
SE-1	Security & Access Control	Access governance (RBAC)	Only authorized roles can access model, data, and configs.	1) Define roles and least-privilege access. 2) Enforce MFA where applicable. 3) Review access quarterly. 4) Maintain audit logs.	RBAC policy; last access review record; sample audit log excerpt.	Policy + logs	NIST			
SE-2	Security & Access Control	Data confidentiality (PHI)	Sensitive data is encrypted and protected end-to-end.	1) Encrypt data at rest and in transit. 2) Tokenize or de-identify where possible. 3) Restrict data egress. 4) Document key management.	Security architecture diagram; encryption and key mgmt description.	Architecture (PDF)	HIPAA, NIST			
SE-3	Security & Access Control	Adversarial robustness testing	You test for and mitigate adversarial/prompt attacks.	1) Identify threat model (e.g., adversarial inputs, data poisoning, prompt injection). 2) Run red-team tests. 3) Implement mitigations/filters. 4) Re-test and document residual risk.	Adversarial test report with scenarios, findings, mitigations.	Test report (PDF)	NIST, CHAI			
SE-4	Security & Access Control	Model integrity & audit logging	Deployed model is tamper-evident and traceable.	1) Sign or checksum model artifacts. 2) Log model version per request. 3) Monitor for anomalies. 4) Preserve logs per retention policy.	Integrity controls description; signed artifact example; log sample.	Control description + logs	NIST			
TT-1	Transparency & Traceability	Explainability for users	Users can understand the basis of outputs.	1) Provide feature importance or saliency (as applicable). 2) Offer plain-language rationale. 3) Document when explanations are less reliable. 4) Include guidance in user docs.	Example explanation output; excerpt from user guide.	Example + user guide	CHAI, NIST			
TT-2	Transparency & Traceability	Model card available	You maintain standardized model documentation.	1) Create model card covering use, data, metrics, limits. 2) Include subgroup results and known risks. 3) Keep current with model versions. 4) Share with customers.	Current model card (with version/date).	Model card (PDF)	CHAI, NIST			
TT-3	Transparency & Traceability	Audit trail of changes	You can trace what changed, when, and why.	1) Track model/data/config versions. 2) Record approvals and rationale. 3) Link to validation results for each change. 4) Retain history per policy.	Version history export; sample change record mapping to results.	Change log (CSV/PDF)	NIST			
TT-4	Transparency & Traceability	Limitations disclosed	Users know when not to rely on the AI.	1) List known weak scenarios. 2) Warn in UI/docs. 3) Prevent out-of-scope use where possible. 4) Review disclosures each release.	Limitations section from docs/UI screenshots; review log.	Docs + screenshots	CHAI			
UX-1	User Experience & Workflow	Workflow integration	The AI fits into existing systems and steps.	1) Integrate with EHR/PACS/ops tools as relevant. 2) Avoid duplicate data entry. 3) Provide APIs/HL7/FHIR as needed. 4) Validate in pilot.	Integration diagram; interface list (APIs/standards); pilot sign-off.	Diagram + sign-off	CHAI			
UX-2	User Experience & Workflow	Training & support provided	Users are trained to use and interpret the AI.	1) Deliver role-based training. 2) Explain limitations and overrides. 3) Provide support channels/SLAs. 4) Track attendance/usage of materials.	Training deck/manual; training roster; support SLA.	Training pack (PDF)	CHAI			
UX-3	User Experience & Workflow	User override & feedback	Users can override and provide feedback.	1) Enable easy override without penalty. 2) Capture feedback in-product. 3) Triage and act on feedback. 4) Communicate changes to users.	UI screenshot of override; feedback log and action tracker.	Screenshots + log	CHAI			
ET-1	Ethical Risk & Societal Impact	Patient disclosure when AI is used	Patients are informed when AI influences care decisions.	1) Draft clear disclosure text. 2) Add to consent/intake as needed. 3) Train staff on usage. 4) Keep records of disclosure use.	Disclosure statement; sample consent; staff comms.	Disclosure pack	CHAI			
ET-2	Ethical Risk & Societal Impact	Human oversight maintained	Humans remain the final decision-makers for critical outcomes.	1) Define decision boundaries. 2) Require human confirmation for critical actions. 3) Log overrides/confirms. 4) Review patterns for overreliance.	Oversight policy; UI/logic showing confirmation; override log sample.	Policy + screenshots	CHAI			
ET-3	Ethical Risk & Societal Impact	Equity & access plan	You consider and mitigate differential access impacts.	1) Identify affected populations/settings. 2) Plan for accessible deployment (languages, resources). 3) Monitor impact by site/population. 4) Adjust rollouts accordingly.	Equity plan with actions/metrics; monitoring snapshot.	Plan (PDF)	CHAI			

ET-4	Ethical Risk & Societal Impact	Ethical review performed	An ethics review has evaluated risks and mitigations.	1) Submit to internal/external ethics board. 2) Address feedback. 3) Record decision and conditions. 4) Re-review on major changes.	Ethics review report/approval; responses to findings.	Review report (PDF)	CHAI			