METHODS & EVIDENCE

The logic behind the conversation.

Risk Continuum CKM

Clinical logic, Level rules, RSS weights, simulator methods, engineering and dated evaluation evidence.

Document version 1.2 · November 5, 2026
Implementation reference: 2d6923ddd2619e4f55bb8791f1306fd60646ed5e
RSS methodology: rss-methodology-v2
Source-specific evidence and clinical evaluation priorities are reported separately.

This appendix brings the clinical pathways, complete Level rules, RSS weights, simulator assumptions and verification records together for review. It includes the current Intermediate-risk minimum of Level 3A and the evidence supporting the implementation.

riskcontinuum.health/methods.html

CONTENTS

01 One assessment, distinct kinds of information.

02 How the Level pathway is assigned.

03 A visible tally of selected findings.

04 Each scenario states how it was modeled.

05 A shared engine, with clinician review in the loop.

06 Evidence for a reproducible clinical workflow.

07 Built for clinician-led prevention care.

08 Published foundations. Explicit implementation choices.

One assessment, distinct kinds of information.

RCCKM combines published risk equations and clinical frameworks with explicit rules for organizing a patient's findings. PREVENT estimates probability; CKM and KDIGO describe clinical context; RCCKM Level and the Risk Signal Score summarize patterns for review.

Meaning and authority of each output

Output	Question it addresses	Basis
PREVENT	What is the modeled probability of a cardiovascular event over a defined horizon?	Published AHA equations, with model eligibility and input-domain checks. [1]
AHA CKM stage	Where do cardiovascular, kidney and metabolic findings sit in the CKM framework?	AHA's disease and risk staging framework, implemented separately from RCCKM Level. [2]
KDIGO G/A context	What do kidney function and albuminuria show, and is chronic kidney disease established?	KDIGO classification plus explicit chronicity and certainty handling. [3]
RCCKM Level	Which prespecified clinical pattern best describes the reviewed findings?	RCCKM's ordered clinical pathways, with the reason for each assignment shown.
Risk Signal Score	Which selected findings contribute to the visible burden summary?	Prespecified points and domain caps, with individual contributors shown for review.

How source evidence becomes a platform rule

Published sources supply the equations, disease definitions and treatment evidence. RCCKM specifies how those inputs are reconciled, which clinical pathway takes priority, and how results are presented together. The Level hierarchy, RSS weights and combined simulator model are documented implementation choices.

Each output retains its own role. Domain-specific recommendations follow the applicable clinical findings and treatment pathways. The assessment presents those reasons alongside the summary measures.

Interpreting the findings together

Measured plaque, actionable drivers and modeled event probability provide complementary information. A patient with a small positive CAC result and a patient with diabetes, severe lipid elevation and reduced kidney function may call for different conversations and management pathways. The assessment keeps the underlying findings available alongside its summaries.

Input meaning is preserved throughout: measured CAC 0, unmeasured CAC, a bounded UACR result and a finding needing confirmation remain distinct. Clinicians review the supplied information and its dates before using the assessment.

How the Level pathway is assigned.

The engine checks clinical disease and coronary calcium first, then evaluates severe lipid states, kidney findings, modeled risk and specified combinations of risk drivers. The first matching base rule establishes the pathway. Intermediate 10-year PREVENT ASCVD risk has a minimum Level of 3A; the separate 3C modifier identifies qualifying combinations within base 3B pathways.

Level names and principal assignment patterns

Level	Principal pattern in the reviewed implementation
1No signals found	No earlier qualifying rule or counted mild signal matches, and the classifier's blocking PREVENT inputs are not missing. Read this in the context of the data available.
2AOne early signal	One counted mild signal, or an isolated kidney finding requiring diagnostic confirmation after higher-priority rules have been considered.
2BSeveral early signals	At least two counted mild signals without an earlier actionable, disease or plaque pathway.
3AModeled risk or single driver	Intermediate 10-year PREVENT ASCVD risk (5-<10%) when a stronger pathway does not apply; qualifying long-term risk; or a single-driver pathway such as LDL-C 160-189 mg/dL, ApoB ≥120 mg/dL or isolated current smoking. The rule and explanation identify the specific pathway.
3BActionable drivers	A qualifying severe-lipid or clinical pathway, or a specified combination of atherogenic, metabolic, kidney, enhancer and/or PREVENT findings, before the 3C modifier.
3CCompounding drivers	A base 3B classification plus any of: 10-year PREVENT ASCVD ≥5%, established diabetes, unconfirmed diabetes-range A1c, or a qualifying kidney finding.
4Plaque present	Quantitative CAC >0 and <300, after clinical disease exclusions.
5Advanced plaque or disease	Clinical ASCVD, established heart failure, or quantitative CAC ≥300. The underlying disease context remains explicit; heart failure alone does not imply an ASCVD diagnosis or ASCVD secondary-prevention lipid targets.
Incomplete	No earlier qualifying rule matches, PREVENT ASCVD is unavailable, and at least one of systolic BP, total cholesterol, HDL-C, eGFR or BMI is recorded as a blocking missing input.

Read the Level with its reason. The code summarizes a clinical pattern, while PREVENT supplies the modeled event probability. The rule identifier and contributing findings explain why the patient follows

that pathway.

Information that keeps its meaning

Numeric CAC determines the plaque-based Level 4/5 pathways. Qualitative incidental calcium remains visible in plaque context, RSS and applicable domain actions. Severe-hypercholesterolemia and suspected FH pathways retain their priority with CAC 0. Kidney measurement categories are interpreted alongside the evidence for chronicity.

Technical appendix A · Complete base-rule sequence and signal definitions

The table lists the base classifier in evaluation order. After that pass, the Intermediate-risk minimum applies to base Levels 1, 2A or 2B; existing 3A and stronger pathways are retained. The 3C modifier then applies to qualifying base 3B pathways. Signal counters and RSS domain caps serve separate purposes.

Ordered base rules, implementation reference 2d6923d

Order	Condition	Base output and rule identifier
1	Clinical ASCVD.	5 LEVEL_5_CLINICAL_ASCVD
2	Established heart failure, HFpEF or HFrEF.	5 LEVEL_5_HEART_FAILURE
3	Quantitative CAC ≥300.	5 LEVEL_5_CAC_GE_300
4	Quantitative CAC >0.	4 LEVEL_4_CAC_POSITIVE
5	Canonical severe hypercholesterolemia (valid LDL-C ≥190 mg/dL or ApoB ≥140 mg/dL), or suspected FH/HeFH.	3B LEVEL_3B_SEVERE_ HYPERCHOLESTEROLEMIA_FH
6	Current kidney category G4/G5 plus A3; compatibility path: eGFR <30 and UACR ≥300.	3B LEVEL_3B_ADVANCED_KIDNEY_ ALBUMINURIA
7	Albuminuria plus at least two actionable drivers or at least one mild signal.	3B LEVEL_3B_ALBUMINURIA_PLUS_ BURDEN
8a	Age 30-59 and 30-year PREVENT ASCVD ≥10%, with HIGH 10-year PREVENT ASCVD category.	3B LEVEL_3B_HIGH_10Y_WITH_ ELEVATED_30Y

Ordered base rules, implementation reference 2d6923d

Order	Condition	Base output and rule identifier
8b	The same elevated 30-year condition with any actionable driver or at least two mild signals. ApoB ≥120 plus premature family history receives a specific reason identifier.	3B LEVEL_3B_30Y_PREVENT_APOB_FAMILY_HISTORY LEVEL_3B_30Y_PREVENT_PLUS_BIOLOGIC_BURDEN
8c	The same elevated 30-year condition without the preceding additional burden.	3A LEVEL_3A_30Y_PREVENT
9	Age 40-59, CAC unmeasured, low short-term PREVENT (<3% or LOW category), ApoB ≥120 or LDL-C 150-<190, and at least two hidden-risk enhancers defined below.	3B LEVEL_3B_HIDDEN_RISK_CLUSTER
10	Age 40-59, CAC unmeasured, low short-term PREVENT, ApoB ≥100, premature family history, and at least two of: early glycemic risk, TG ≥150, valid LDL-C ≥130, or CKM context.	3B LEVEL_3B_APOB_FAMILY_METABOLIC
11	Valid LDL-C 160-189.999 mg/dL. At least two actionable drivers or two mild signals gives 3B; otherwise 3A. The young family/metabolic helper changes the explanation within the 3B branch.	3B or 3A LEVEL_3B_LDL_160_189_CONVERGENCE LEVEL_3A_LDL_160_189
12	Rheumatoid arthritis plus premature family history, 10-year ASCVD <3%, 30-year ASCVD unavailable or <15%, without established diabetes, a kidney signal or severe hypercholesterolemia.	3A LEVEL_3A_INFLAMMATORY_CONTEXT
13	At least two enhancer-burden flags plus at least one metabolic-burden flag.	3B LEVEL_3B_MULTIPLE_ENHANCERS
14	Elevated Lp(a) plus at least two metabolic-burden flags.	3B LEVEL_3B_LPA_METABOLIC_BURDEN
15	Current smoking: established diabetes or at least two actionable drivers gives 3B; otherwise 3A.	3B or 3A LEVEL_3B_SMOKING_PLUS_BURDEN LEVEL_3A_SMOKING
16	BORDERLINE, INTERMEDIATE or HIGH 10-year PREVENT ASCVD category, plus at least two actionable drivers, albuminuria, or at least three mild signals.	3B LEVEL_3B_PREVENT_PLUS_BURDEN
17	HIGH PREVENT category, established diabetes, or a clinical kidney pathway (established CKD or current G3a-G5). Compatibility path uses a qualifying kidney signal.	3B LEVEL_3B_ACTIONABLE_CLINICAL_PATHWAY

Ordered base rules, implementation reference 2d6923d

Order	Condition	Base output and rule identifier
18	Kidney confirmation pathway: current A2/A3 or documented albuminuria without an earlier clinical kidney pathway.	2A LEVEL_2A_ISOLATED_KIDNEY_CONFIRMATION
19	ApoB ≥120 mg/dL.	3A LEVEL_3A_APOB_GE_120
20	At least two counted mild signals.	2B LEVEL_2B_MULTIPLE_MILD_SIGNALS
21	Exactly one counted mild signal.	2A LEVEL_2A_SINGLE_MILD_SIGNAL
22	PREVENT ASCVD unavailable and one or more listed blocking inputs missing: systolic BP, total cholesterol, HDL-C, eGFR or BMI.	Incomplete LEVEL_INCOMPLETE_MISSING_PREVENT_INPUTS
23	No earlier rule matches.	1 LEVEL_1_MINIMAL_RISK

The Intermediate-risk minimum

An authoritative Intermediate 10-year PREVENT ASCVD category (5-<10%) raises a base Level 1, 2A or 2B to 3A under LEVEL_3A_INTERMEDIATE_PREVENT. Existing 3A, 3B/3C, plaque and clinical-disease pathways retain their assignments. An incomplete assessment retains its missing-input status. Treatment recommendations continue to follow their own clinical pathways.

The 3C modifier

After any base 3B rule, the engine checks 10-year PREVENT ASCVD ≥5%, established diabetes or unconfirmed diabetes-range A1c, and the kidney-signal helper. Any positive condition changes the displayed Level to 3C, retains the base rule identifier, and records the convergence drivers separately. RSS is not an input to this modifier.

Definitions used by the counters

• **Early glycemic risk:** the canonical glycemia interpretation distinguishes prediabetes, an unconfirmed diagnostic-range A1c and established diabetes. Diagnosis confirmation is handled separately from the risk signal.

• **CKM context:** BMI ≥30 kg/m², OSA or MASLD, counted together as one context flag.

• **BP signal:** represented BP treatment, hypertension/elevated-BP state, or at least two usable BP pairs with systolic ≥130 or diastolic ≥80 mmHg. The Level helper's BP counter is distinct from RSS's

disease-authority requirement.

• **Elevated Lp(a):** ≥125 nmol/L or ≥50 mg/dL in the reported unit. No fixed conversion between units is applied.

• **Kidney signal:** established CKD, current G3a-G5, current A2/A3, or supported problem-list albuminuria. The compatibility path uses eGFR <60, UACR ≥30, or documented albuminuria. Current abnormal categories can trigger a signal while chronicity remains unconfirmed.

• **Counted mild signals:** early glycemic risk; TG ≥150; ApoB 80-<120, or, when ApoB is absent, atherogenic findings (valid LDL-C ≥100, non-HDL-C ≥130 or TG ≥150); CKM context; elevated Lp(a); premature family history; reproductive history; inflammatory disease; HIV; South Asian ancestry; Filipino ancestry; breast arterial calcification; hsCRP ≥2; and BORDERLINE or INTERMEDIATE PREVENT ASCVD category. The lipid fallback is one flag and can overlap with the separate TG flag.

• **Actionable drivers:** albuminuria, reduced current G category, established diabetes, smoking, LDL-C 160-189.999, severe hypercholesterolemia, suspected FH, BP signal, ApoB ≥100 (one ApoB driver, with a distinct label at ≥120), elevated Lp(a), premature family history, reproductive history, early glycemic risk, TG ≥150 and CKM context.

• **Metabolic-burden counter:** one flag each for early glycemic risk, established diabetes, BP signal, CKM context, TG ≥150, ApoB ≥100, non-HDL-C ≥160 and valid LDL-C ≥130. These flags can describe related biology.

• **Enhancer-burden counter:** elevated Lp(a), premature family history, any reproductive flag, inflammatory disease, HIV, active cancer, cancer survivorship with represented life expectancy >2 years, South Asian ancestry, Filipino ancestry and breast arterial calcification. Each is a separate Boolean flag in this implementation.

• **Hidden-risk enhancer counter:** elevated Lp(a), premature family history, inflammatory disease, South Asian ancestry, Filipino ancestry, gestational diabetes, preeclampsia, early/premature menopause (one flag) and breast arterial calcification.

• **Inflammatory signal:** a generic inflammatory-disease flag, rheumatoid arthritis, SLE, psoriasis, inflammatory arthritis or IBD. **Reproductive flags:** early/premature menopause, preeclampsia, gestational hypertension/diabetes, preterm delivery, small-for-gestational-age infant, recurrent pregnancy loss, PCOS/irregular menses and early menarche.

10-year PREVENT ASCVD categories are Low <3%, Borderline 3-<5%, Intermediate 5-<10%, and High ≥10%. At all three boundaries, the classifier uses unrounded calculated risk when matching model provenance is available. Reported values remain rounded; clinician-provided and legacy values retain their documented authority.

Boundary example. A calculated result of 9.996% remains Intermediate even if its reported numeric value rounds to 10.00%. If no stronger pathway applies, the Level is 3A. An exact 10% result is High. These category boundaries use the same precision policy; the minimum Level adds no extra tolerance band.

A visible tally of selected findings.

RSS v2 is a transparent tally of selected findings. Plaque receives greater weight, major biologic states receive intermediate weight, and contextual findings receive smaller weights. The weights and overlap rules below make each contribution visible.

The total is the sum of displayed positive contributions, with domain caps applied where specified. Findings can remain visible as context after a cap is reached. Missing measurements remain unknown and contribute no points; interpret the total alongside the available inputs.

How overlap is handled

• **Lipids:** ApoB, or LDL-C ≥190 when ApoB is absent, is allocated first; suspected FH next; triglycerides last. These contributions share a 12-point cap. Lp(a) is a separate measured domain.

• **Plaque:** a numeric CAC result supersedes qualitative incidental CAC. The two do not stack.

• **BP:** disease authority or repeated credible elevations is required. Diagnosis, therapy and current measurement produce one contribution, up to 8 points. An isolated reading alone does not qualify.

• **Inflammatory disease and ancestry:** multiple supported labels within each domain contribute once. HIV remains a separate domain; hsCRP has its own measured contribution.

• **Reproductive history:** 2 points per supported marker, up to 4. Early and premature menopause are reconciled into one menopause item.

• **Kidney findings:** eGFR and albuminuria are additive, reflecting separate G and A axes. Their contributions do not by themselves establish CKD chronicity.

Technical appendix B · Every RSS band and point value

These tables are generated from the reviewed engine's canonical RSS methodology export. The downloadable JSON also supplies exact continuous interval endpoints and source fingerprints. Integer-style band labels below follow the engine's display convention.

CAC

CAC: prespecified RSS weights

Finding	Points	Scoring note
1-99	10	CAC 1-99 indicates mild plaque burden.
100-299	20	CAC 100-299 indicates moderate plaque burden.
300-999	30	CAC 300-999 indicates severe plaque burden.
>=1000	40	CAC >=1000 indicates extreme plaque burden.

ApoB

ApoB: prespecified RSS weights

Finding	Points	Scoring note
80-99 mg/dL	4	ApoB 80-99 represents borderline elevation.
100-129 mg/dL	8	ApoB 100-129 represents elevated ApoB.
>=130 mg/dL	12	ApoB >=130 is a higher RCCKM RSS particle-burden scoring band, not a treatment threshold.

LDL-C fallback

LDL-C fallback: prespecified RSS weights

Finding	Points	Scoring note
LDL-C >=190 mg/dL when ApoB unavailable	10	LDL-C >=190 mg/dL indicates severe hypercholesterolemia.

Lp(a)

Lp(a): prespecified RSS weights

Finding	Points	Scoring note
75-124 nmol/L / 30-49 mg/dL	2	Lp(a) 75-124 nmol/L is the consensus grey zone and a small risk-refinement signal.
125-249 nmol/L / 50-99 mg/dL	8	Lp(a) 125-249 nmol/L indicates elevated Lp(a).
250-429 nmol/L / 100-179 mg/dL	12	Lp(a) 250-429 nmol/L indicates higher elevated Lp(a).
>=430 nmol/L / >=180 mg/dL	15	Lp(a) >=430 nmol/L indicates very high Lp(a) elevation.

hsCRP / inflammatory risk

hsCRP / inflammatory risk: prespecified RSS weights

Finding	Points	Scoring note
Single hsCRP >=2 mg/L without established persistence/context	2	Unconfirmed measurement; confirm persistence or clinical context.
2-4.9 mg/L	4	hsCRP 2-4.9 mg/L indicates elevated inflammatory risk.
>=5 mg/L	7	hsCRP >=5 mg/L indicates high inflammatory risk.

Blood pressure / hypertension

Blood pressure / hypertension: prespecified RSS weights

Finding	Points	Scoring note
Established/treated and controlled or current BP unavailable	3	Underlying hypertension burden persists despite control; an isolated reading alone does not qualify.
Established/repeated stage 1 BP	4	Mild established BP burden.
Established/repeated stage 2 BP	6	Major established BP burden.
Repeated markedly/severely elevated BP	8	Maximum BP-domain burden; acute action remains separate from RSS.

Kidney disease

Kidney disease: prespecified RSS weights

Finding	Points	Scoring note
45-<60	5	eGFR 45-<60 indicates mild kidney function reduction.
30-<45	8	eGFR 30-<45 indicates moderate kidney function reduction.
15-<30	12	eGFR 15-<30 indicates severe kidney function reduction.
<15	15	eGFR <15 indicates kidney failure risk.
30-<300 mg/g	6	UACR 30-<300 indicates moderately increased albuminuria.
>=300 mg/g	10	UACR >=300 indicates severely increased albuminuria.

Triglycerides

Triglycerides: prespecified RSS weights

Finding	Points	Scoring note
150-499 mg/dL	4	Triglycerides 150-499 indicates hypertriglyceridemia.
>=500 mg/dL	4	Triglycerides >=500 indicates severe hypertriglyceridemia; pancreatitis urgency is handled outside RSS.

Diabetes / dysglycemia

Diabetes / dysglycemia: prespecified RSS weights

Finding	Points	Scoring note
5.7-6.4%	2	A1c 5.7-6.4% indicates prediabetes-range dysglycemia.
Diabetes / diabetes-range glycemia	8	Diabetes is a risk enhancer for ASCVD.

Smoking and family history

Smoking and family history: prespecified RSS weights

Finding	Points	Scoring note
Current smoking	8	Current smoking is a modifiable risk factor.
Premature family history	3	Premature first-degree family history is a risk enhancer.
Suspected FH / HeFH	8	Up to 8 points, limited by the 12-point atherogenic-lipid overlap cap.

Incidental CAC

Incidental CAC: prespecified RSS weights

Finding	Points	Scoring note
Present	10	Incidental coronary calcium is qualitative plaque evidence.
Severe	20	Severe incidental coronary calcium is qualitative high-burden plaque evidence.

Contextual enhancers

Contextual enhancers: prespecified RSS weights

Finding	Points	Scoring note
Chronic inflammatory disease, HIV, OSA, MASLD, or higher-risk ancestry	2	Each distinct contextual domain is scored once.
Reproductive risk markers	2 each; maximum 4	The reproductive-history domain is capped so correlated markers do not stack without limit.

Interpretation details

For BP, the implementation classifies stage 1 at systolic ≥130 or diastolic ≥80, stage 2 at ≥140 or ≥90, markedly elevated at ≥160 or ≥100, and severely elevated at ≥180 or ≥120 mmHg, with higher categories evaluated first. One severe current reading in established hypertension receives 6 points; 8 requires at least two markedly/severely elevated historical readings. Below-stage-1 or unavailable current BP with established disease receives 3. Clinical urgency is evaluated outside RSS.

A single hsCRP ≥2 mg/L without persistence or supportive chronic context receives 2 points. The 4- and 7-point bands require the corresponding persistence/context authority. Established diabetes or A1c ≥6.5% receives the 8-point glycemic contribution; an isolated diagnostic-range result remains subject to separate diagnosis-confirmation rules.

The contextual ancestry domain counts supported South Asian and/or Filipino flags once. These are RSS context weights; PREVENT uses its published predictor set. Clinical evaluation will examine the calibration, fairness and incremental usefulness of these contextual contributions.

Reproductive items include supported menopause and menarche age/flags, preeclampsia, gestational hypertension/diabetes, preterm delivery, small-for-gestational-age infant, recurrent pregnancy loss and PCOS/irregular menses. A reproductive item adds 2 points only within the 4-point cap; all represented items remain available as detail.

RSS is read alongside its individual contributors, PREVENT probability and established disease context. Its points are a burden summary rather than a percentage out of 100.

Each scenario states how it was modeled.

Each plan forks the reviewed worksheet and reruns the clinical engine. The Level, RSS and domain interpretation reflect that scenario. The ASCVD probability comparison uses two explicitly distinguished methods.

Trial-derived LDL-C and systolic BP scenarios

For LDL-C lowering, the simulator uses the CTT 2010 relative-effect estimate of 0.78 per 1 mmol/L (38.67 mg/dL) reduction. For systolic BP lowering, it uses Ettehad 2016's estimate of 0.80 per 10 mmHg reduction. [4] [5] The engine scales each factor to the selected change, multiplies the factors, and applies them to baseline survival:

f_lipid = 0.78 ^ ((LDL_before - LDL_plan) / 38.67)

f_BP = 0.80 ^ ((SBP_before - SBP_plan) / 10)

f_combined = product of applicable factors

p_plan = 1 - (1 - p_reference) ^ f_combined

Here, p_reference is a proportion between 0 and 1. For lipid/BP changes alone, it is the baseline PREVENT ASCVD estimate. In a mixed scenario, supported risk-factor changes first recalculate PREVENT, then the lipid/BP factors are applied. Each factor is rounded to three decimals before multiplication; risk percentages are reported to two decimal places. Display precision is separate from the uncertainty of a modeled projection. Changes smaller than 0.5 mg/dL LDL-C or 0.5 mmHg systolic BP add no trial-effect factor.

The combined projection assumes proportional effects and multiplicative combination of the lipid and BP factors. CTT studied major vascular events, including revascularization; Ettehad studied major cardiovascular outcomes, including heart failure. Their endpoints and populations differ from PREVENT ASCVD. The simulator applies these published relative effects to the selected PREVENT horizon as a prespecified modeling assumption.

Recalculated risk-factor scenarios

Supported smoking, HbA1c and UACR changes rerun the applicable PREVENT model. A new value affects probability when that predictor belongs to the baseline model. Comparisons preserve the baseline model and diabetes status; changes outside that comparison receive an explanation. These scenarios represent changes in the modeled risk profile, rather than the full benefit of a particular drug.

Prespecified treatment assumptions

Starting a moderate-intensity statin assumes LDL-C and non-HDL-C decrease by 35%, and ApoB by 30%, with HDL-C held constant for the non-HDL calculation. Intensification starts from the **current on-treatment lipid value**. For a moderate-intensity regimen, the LDL/non-HDL multiplier is 0.50/0.65, approximately 0.769: the assumed remaining fractions of an untreated value on high- and moderate-intensity therapy. A current LDL-C of 100 mg/dL therefore becomes approximately 77 mg/dL, an additional reduction of about 23%. ApoB is multiplied by 0.80. For a represented low-intensity regimen, the current LDL/non-HDL value is multiplied by 0.50/0.80 (0.625), and current ApoB by 0.70. Lipid values are rounded as implemented. These are scenario defaults; observed response requires follow-up.

If LDL-C is absent, the selected total-cholesterol change, with baseline HDL-C held constant, can serve as a labeled non-HDL-C proxy for the lipid effect. ApoB affects interpretation but receives no additional trial ASCVD reduction. Treatment effects and direct predictor changes are separated to avoid counting the same lipid/BP change twice.

Horizon and scope

RCCKM reports supported 10-year PREVENT estimates at ages 30-79 and 30-year estimates at ages 30-59. Its eligibility policy directs established clinical cardiovascular disease, established heart failure and CAC ≥300 to their clinical pathways. Eligibility and input-domain checks follow the implemented source guidance. [1]

Scenario estimates cover ASCVD over the supported horizon. Drug harms, life expectancy, annual event timing and statin withdrawal are outside these calculations. A 30-year trial-effect projection assumes sustained benefit beyond the roughly five-year statin trial follow-up; connecting lines are visual guides. Adherence, competing risks, changing health and subsequent treatment remain part of the clinical discussion.

A shared engine, with clinician review in the loop.

1. **1 · Reconcile the source**A rule-based parser proposes structured values, dates, histories and medication states from supported chart text. Missing fields and conflicts remain visible.

2. **2 · Confirm the worksheet**The clinician reviews or edits the proposed inputs. Structured disease, measurement, treatment and certainty states carry their own meanings.

3. **3 · Run the clinical engine**The Python engine evaluates inputs, published calculations and explicit domain rules, producing structured results and traceable reasons.

4. **4 · Review and communicate**The interface presents the assessment, scenarios, clinician note and patient handout. The clinician reviews these outputs and makes the clinical decisions.

Deterministic core and optional AI

Confirmed inputs, the same engine/evidence version and the same evaluation context follow the same deterministic calculation and decision logic. Next.js, FastAPI and the retained Streamlit interface consume the shared engine's results. Renderers select and format structured clinical facts rather than independently recomputing them.

Optional AI commentary is labeled separately and presented for clinician review. It can vary between requests. Calculated risks, Level, RSS and staging come from the shared deterministic engine; the clinician reviews the note, diagnosis support and patient handout before use.

Input meaning and assessment lifecycle

The engine distinguishes absent results from measured zero, quantitative results from bounded/nonnumeric results, current treatment from documented non-use, and established disease from a finding that needs confirmation. Clinical context and observation dates matter. Input changes invalidate a stale assessment and dependent outputs, prompting reassessment.

Kidney interpretation retains G/A measurement categories separately from CKD certainty. Glycemia separates established diabetes from unconfirmed diagnostic-range measurements. BP interpretation separates a usable reading, disease authority and treatment context. These distinctions let an abnormal result prompt review without silently creating a confirmed diagnosis.

Versioning and interoperability

Result metadata includes application version, build identifier, rule-set version, evidence-map hash and generation timestamp. Rule traces expose triggering inputs and rule identifiers where implemented. A reproducible review must specify the software/evidence version, confirmed inputs, evaluation date and relevant configuration.

The public workflow supports synthetic or de-identified text/manual entry and clinician-controlled output transfer. FHIR R4 normalization, SMART sandbox and CDS Hooks adapters provide integration capabilities. Production EHR integration requires organization-specific configuration and approval.

Evidence for a reproducible clinical workflow.

Verification checks the calculations, rule implementation and path from source information to the assessment and copied output. Synthetic cases make threshold behavior, missing-data handling and regressions reproducible across software versions.

Evidence register, checked November 5, 2026. Test dates are Eastern Time.

Verification layer	Source and date	Completed checks	What it covers
Routine software suite	789d382 November 4, 2026	**7,277 passed; 160 skipped** on each of Linux and Windows.	Parser, clinical rules, invariants, API contracts and output regressions.
Harness and browser regression checks	789d382 November 4, 2026	**379 passed; 15 skipped.** The separate Next.js build and browser-regression job also passed.	Harness integrity, worksheet behavior, rendering, copied output and supported web workflows.
Provider-free deterministic replay	789d382 November 4, 2026	**600 journeys passed**: 500 primary synthetic records plus 100 format variants. Approved-baseline comparison: **zero changed journeys**.	Repeatable parsing, worksheet construction, engine outputs, reconciliation and report contracts.
FHIR/paste parity	789d382 November 4, 2026	**500 synthetic journeys passed.**	Agreement between supported FHIR normalization and paste-based inputs.
Independent PREVENT reference	Pinned AHA implementation; reused after source/input equivalence checks.	**535 direct matches; 5 governed boundary-projection matches; 60 clinically ineligible journeys.** All 600 records and two published examples reproduced with zero differences.	Numerical inputs, eligibility, model selection and endpoints against the recorded reference.
Browser and AI semantic QA	Earlier revision 933e1f6 October 31, 2026	**600 completed reviews** through the retained Streamlit driver; passing recorded findings.	Historical review of synthetic browser journeys, interpretation and copy. This record stays linked to its tested revision.

The current implementation reference is merged revision `2d6923d`, which incorporates the tested application head `789d382`. The downloadable evidence register records full source identifiers, workflow

links, receipt excerpts and review scope. Supporting source and validation artifacts are available for review.

What the verification layers examine

• **Calculations:** comparison with the pinned independent PREVENT reference, including supported models, eligibility and governed boundary projections.

• **Clinical rules:** synthetic patterns and threshold cases for disease, plaque, lipids, glycemia, kidney findings, BP and missing information.

• **Regressions:** comparison of parsed values, structured results and copied reports against physician-approved expectations.

• **Browser workflows:** ingestion, reconciliation, worksheet state, assessment presentation, accessibility and copied output.

• **Semantic QA:** model-assisted review of synthetic journey evidence against a specified rubric, with findings and adjudications retained by source version.

The next stage: independent clinical evaluation

Current evidence addresses software implementation and synthetic workflows. Independent clinician review and prospective usability studies are needed to evaluate clinical usefulness, decision quality and real-world impact. RCCKM Level and RSS have not been externally calibrated as event-prediction models, and the combined simulator projection has not been prospectively validated.

Planned evaluation priorities

Question	Evaluation approach
Are the interpretations and pathways clinically appropriate?	Independent clinician adjudication of prespecified cases, with agreement, disagreements and safety findings reported.
Does the workflow help a busy clinician?	Task completion, time to decision, usability, workload and confidence in a prospective workflow study.
Does it support the prevention conversation?	Patient understanding, shared decision making and clinician use over time.

Technical appendix C · Source versions and evidence records

Each validation receipt identifies the application revision, numerical reference, corpus and workflow used. The November 4 CI run passed all four jobs: Linux, Windows, Next.js and the Critic harness. The merged implementation adds the consistent 10% precision policy and Intermediate-risk minimum of 3A.

The October 31 browser/AI semantic record belongs to revision 933e1f6e6c1dbbbd5fd1af5ed84566c6ea3cbb45, source SHA-256 ae6c5c7fffdfb6a2d4741ba57a5e68233a9c7a20ae884e3d7a584eebd1bd6af6. The current merge's separate semantic gate was waived by the physician owner; it is recorded as waived rather than as a new semantic pass.

Original findings, adjudications and superseded runs remain in the development record. Software revisions and recorded test sources are listed separately from deployment status. Test counts refer to software checks and synthetic journeys, not clinical study participants.

Built for clinician-led prevention care.

RCCKM supports assessment and discussion alongside clinical judgment. Review imported values, dates, unresolved findings and outputs before use. The care plan also reflects diagnostic confirmation, contraindications, adverse effects, tolerability, treatment burden, cost and patient preferences.

Use synthetic or properly de-identified information in the public demonstration; keep patient identifiers out of that workflow. Optional AI features have separate data-processing terms, described in the current Privacy Policy and Terms.

Published foundations. Explicit implementation choices.

1. American Heart Association. PREVENT equations and clinical resources.
Published equation framework, primary-prevention use and optional predictors. Includes the current clinician FAQ and source-method links.

2. American Heart Association. Cardiovascular-Kidney-Metabolic Health Presidential Advisory. 2023.
Integrated CKM framework and staging.

3. KDIGO. 2024 Clinical Practice Guideline for the Evaluation and Management of CKD.
Kidney function, albuminuria classification and chronicity, implemented from the 2024 guideline.

4. CTT Collaboration. Efficacy and safety of more intensive lowering of LDL cholesterol. Lancet. 2010;376:1670-1681.
Trial-derived LDL-C effect used in scenario modeling. DOI: 10.1016/S0140-6736(10)61350-5.

5. Ettehad D, et al. Blood pressure lowering for prevention of cardiovascular disease and death. Lancet. 2016;387:957-967.
Trial-derived systolic BP effect used in scenario modeling. DOI: 10.1016/S0140-6736(15)01225-8.

6. 2026 ACC/AHA multisociety Guideline on the Management of Dyslipidemia.
Lipid, biomarker and CAC prevention context. Guideline support for a finding does not validate the RSS

weight assigned to it.

7. European Atherosclerosis Society. Lipoprotein(a) consensus statement. 2022.
Unit-specific Lp(a) concentration context; RSS weights remain prespecified platform choices.

8. Budoff MJ, et al. Ten-year association of coronary artery calcium with ASCVD events: MESA. 2018.
Observed CAC risk gradients; the RSS point scale itself is not derived or validated by this study.

References identify underlying scientific work. RCCKM is independently developed by Risk Continuum Health LLC. References do not imply organizational endorsement.

Version 1.2 · November 5, 2026. Updated Level logic, consistent PREVENT boundary handling and the latest completed software checks. Implementation reference: 2d6923ddd2619e4f55bb8791f1306fd60646ed5e.

The downloadable registers provide source fingerprints, rule identifiers, exact RSS intervals and dated verification receipts. Earlier versions, findings and adjudications remain in the development record.