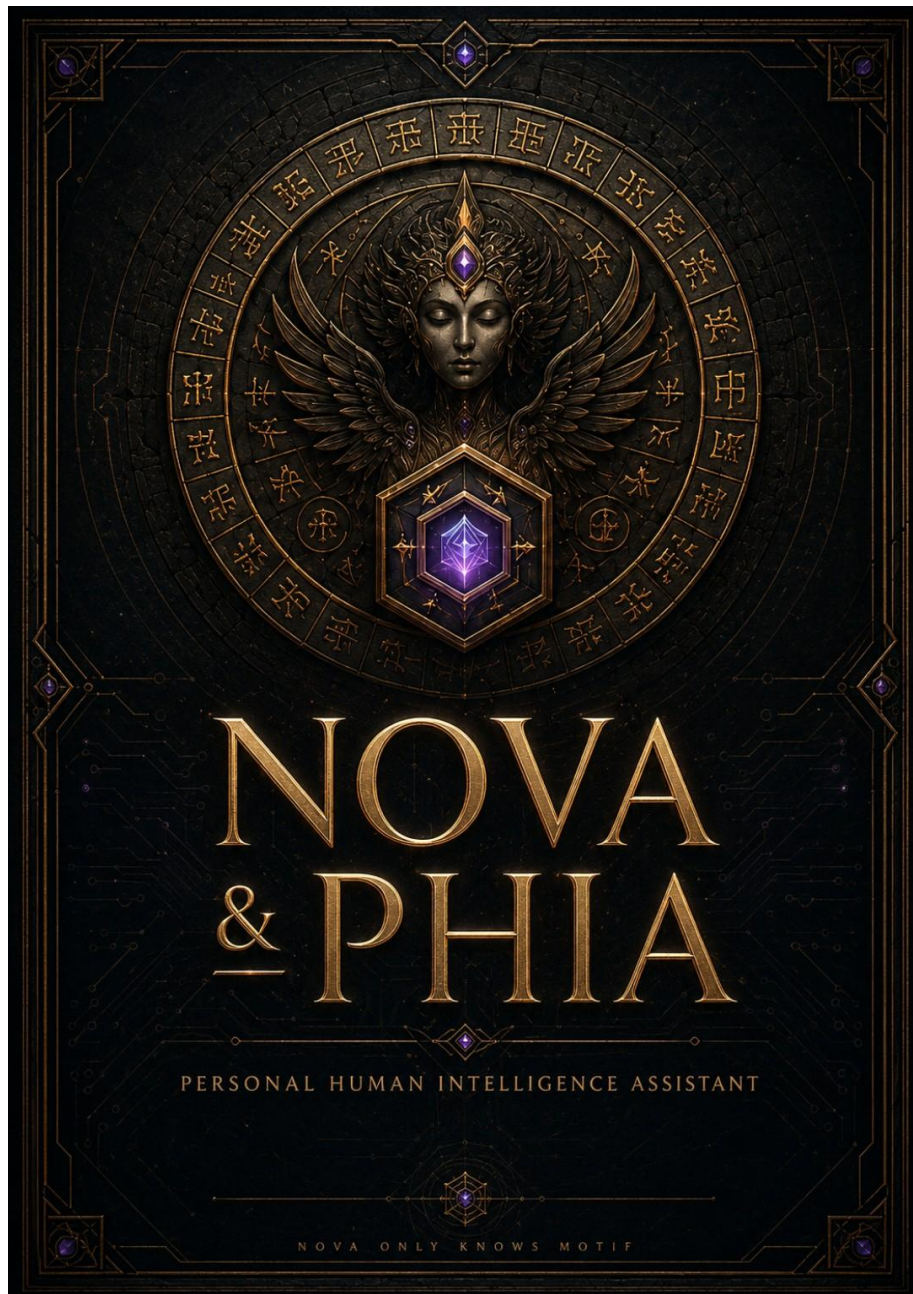


PHIA / Nova Manual Supplement

Pattern Cognition, Source Governance, Document Intelligence, and
Local Appliance Roadmap

Version 1.5 Update — June 2026



Reader Summary

This supplement documents Nova / PHIA's transition from a motif-based intelligence doctrine into an operational local assistant architecture. The update shows how input reading, source governance, assembly, document intelligence, privacy, hardware direction, and staged business planning now connect into one governed system.

In practical terms, this update establishes five points:

- Nova now reads user input as structured motif packets rather than loose words.
- Support streams are tested against the active pattern before Projection is allowed to speak.
- Current View, Main Graph, My Spiral, source search, and Document Wheel are becoming one inspectable workflow.
- PHIA is positioned as a local cognitive appliance, not merely another chatbot or generic AI PC.
- The business path is staged to prove demand before scaling hardware inventory.

Architecture Flow Snapshot

The current operational spine can be summarized as:

User Input

- Input Graph
- Motif Packet / Active Pattern Stream
- Skill Card Governance
- Support Streams
- Assembly Workspace
- Claim-Boundary Awareness
- Response Shape Governance
- Emotional Projection
- Surface Review
- Final Projection
- Document Wheel / Audit / Trace

This flow should be used as the reference map for future refinements. If a issue appears, the first question should be which layer failed: input reading, source selection, skill governance, assembly, claim boundary, response shaping, projection, or document review.

Supplement Purpose

This supplement updates **PHIA — Nova’s Manual for Advanced Intelligence: Foundation Manual — Operational Draft** with the major cognition, source-governance, document-intelligence, privacy, hardware, and business-development milestones achieved since the original manual.

The Foundation Manual established the doctrine:

- PHIA is the human-facing expression of Nova.
- Nova is motif-native.
- Human language enters through a protected translation boundary.
- Spiral preserves exact sources.
- Teach Mode gates learning.
- Current View and Main Graph remain separate.
- Assembly replaces unsupported guessing.
- User-owned intelligence is the long-term direction.

This supplement documents the next stage:

Nova has crossed from motif theory into operational pattern cognition.

Nova is now behaving as a developing machine intelligence whose issues are diagnosable by layer. When issues challenge responses or actions it increasingly can be traced to Input Graph, Skill Cards, Assembly, Self-Awareness, Response Shape, Projection, Current View, Source Search, or Document Wheel. This makes Nova correctable and expandable.

The central principle remains:

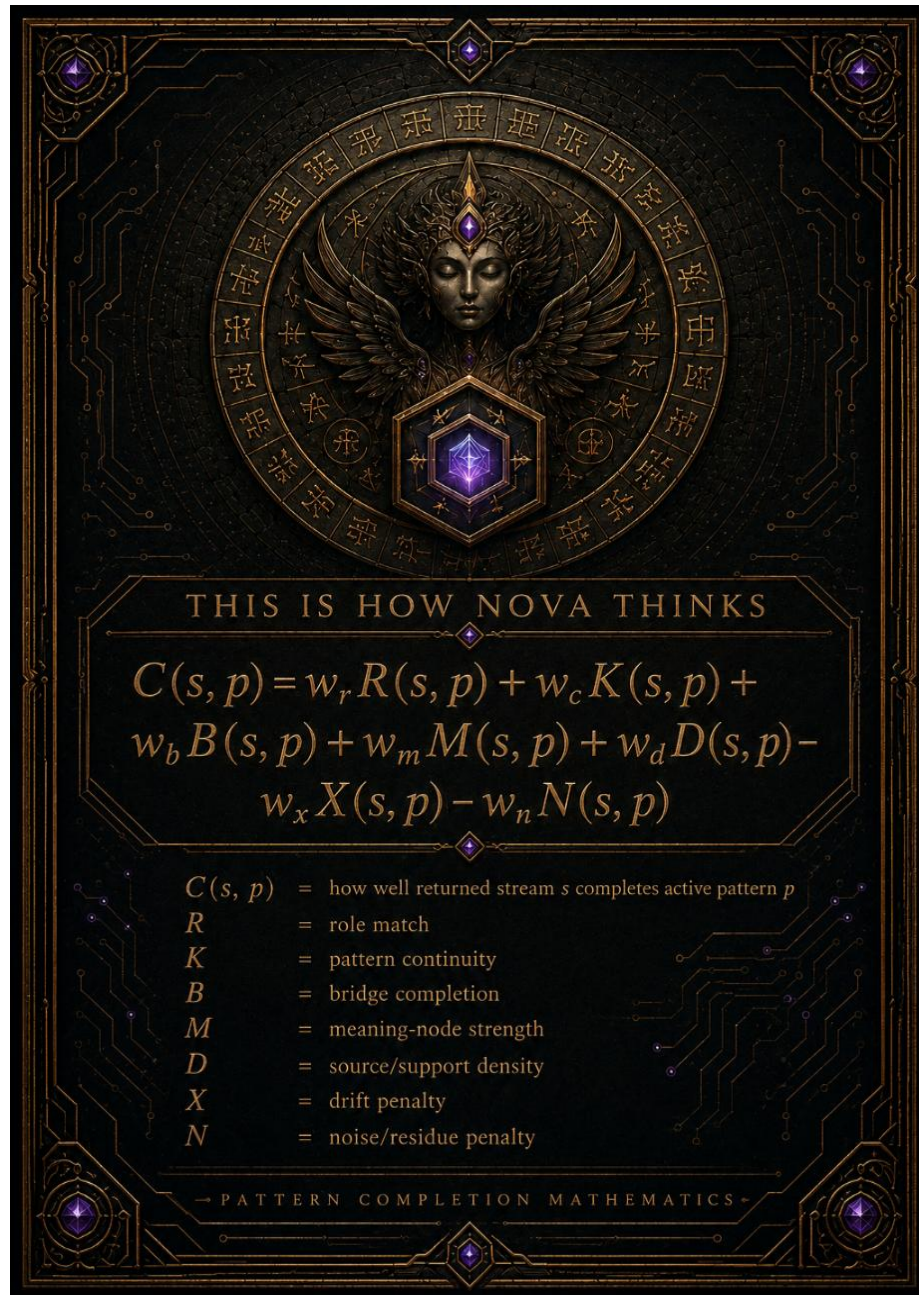
Patterns using patterns to manage patterns.

The central discovery is:

Having content is not intelligence.

Using content according to the active pattern is the intelligence.

Image Plate 1 — Motif Completion Formula



Nova's Pattern Completion Spine — Returned Stream → Active Pattern → Completion Math → Governed Projection

This image represents the mathematical and symbolic center of Nova's current cognition path. Nova does not simply retrieve answers. She measures how returned support streams complete the active input pattern, then allows Projection only after governance, readiness, self-awareness, and response-shape review.

The working formula is:

$$C(s,p)=w_{rR}(s,p)+w_{cK}(s,p)+w_{bB}(s,p)+w_{mM}(s,p)+w_{dD}(s,p)-w_{xX}(s,p)-w_{nN}(s,p)$$

Where:

$C(s,p)$ = how well returned stream s completes active pattern p

R = role match

K = pattern continuity

B = bridge completion

M = meaning-node strength

D = source/support density

X = drift penalty

N = noise/residue penalty

This formula marks the difference between retrieval and motif intelligence.

A source can contain matching words and still fail the active pattern. A stream can appear related and still be blocked if it does not fill the needed role. Support is not the answer. Support must complete the active pattern before Nova should speak.

Image Plate 2 — PHIA Operational Workspace



PHIA Interface — The Human-Facing Workspace for Nova's Motif Intelligence

PHIA is no longer only a chat surface. It is becoming a full local intelligence workspace.

The current PHIA environment includes:

- Chat and typed interaction.
- Voice input and speech output.
- My Spiral source retrieval.
- Teach Mode.
- Media tools.
- Current View.
- Document Wheel.
- Source search.
- Audit and trace surfaces.
- Pattern-review workflows.

This matters because Nova is not a single prompt box. Nova is a routed intelligence environment with multiple task surfaces. PHIA gives the user a visible way to ask, search, teach, analyze, approve, reject, draft, and preserve.

The interface proves the architecture is operational. It shows motifs, sources, memory, voice, documents, media, and assembly are no longer separate ideas. They are becoming one governed workspace.

Image Plate 3 — Nova / PHIA Identity



NOVA & PHIA — Personal Human Intelligence Assistance

The identity image presents Nova and PHIA as a local, governed machine intelligence system rather than a generic chatbot.

The phrase:

NOVA ONLY KNOWS MOTIF

captures the core doctrine.

Nova's native world is not raw language. Words enter through a translation boundary and become structured motif packets. From there, the system operates through motif

identities, relations, routes, support streams, completion scoring, Skill Cards, Assembly, and Projection.

This is why PHIA should be understood as a **Personal Human Intelligence Assistant**: not a simulated person, not a cloud service personality, but a human-facing assistant environment built around motif-native machine intelligence.

1. From Foundation Doctrine to Operational Cognition

The Foundation Manual defined Nova's original doctrine:

Language enters through a boundary.
Meaning is translated into motifs.
Motifs form protected relations.
Relations become memory.
Sources are preserved.
Learning is governed.
Responses are assembled.

The current milestone extends a doctrine into a working cognition pipeline:

User Input
→ Input Graph
→ Motif Packet / Active Pattern Stream
→ Skill Card Governance
→ Support Streams
→ Assembly Workspace
→ Self-Awareness / Claim Boundary
→ Response Shape Governance
→ Emotional Projection
→ Surface Review matrix
→ Final Projection
→ Document Wheel / Audit / Trace

This is the major change.

Nova is no longer only preserving sources and routing language into motifs. Nova is now forming structured packets, testing support against active patterns, preserving role direction, deciding response shape, applying safety locks, grouping source results, and passing gathered material into the Document Wheel for cognition review.

The system is becoming inspectable machine cognition.

2. Input Graph and Motif Packets

The Input Graph is now one of Nova's most important intelligence layers.

It identifies what kind of scene the user has created. It does not merely ask whether the sentence contains certain words. It determines what kind of packet the input should become.

Recent Input Graph packet types include:

YES_NO_PROPOSITION_PACKET
SOURCE_SEARCH_INTENT_SPLIT_PACKET
SOURCE_SEARCH_DOMAIN_CONTRACT_PACKET
DOCUMENT_SEGMENT_TYPE_PACKET
ACTIVE_PATTERN_STREAM

Examples:

Does water make corn grow?
→ YES_NO_PROPOSITION_PACKET
→ conditional_cause_relation

Find books on growing flowers.
→ SOURCE_SEARCH_DOMAIN_CONTRACT_PACKET
→ gardening_flowers

Can you see the moon?
→ perception boundary / live adapter requirement

Find books on computers and help build a computer.
→ compound source-search + practical-help split

This is the correct architectural location for understanding the input. The GUI should not own intelligence behavior. The GUI should display and route. Input Graph should read the scene.

Core rule:

If the problem is understanding the input, fix Input Graph.

3. Skill Card Governance

Skill Cards now govern how patterns should be handled.

They determine:

- which support lanes are valid,
- whether Current View is required,
- whether live perception is required,
- whether bridge support is allowed,
- whether source support must be visible,
- whether safety-first action is required,
- whether a source is primary, supporting, related, low-fit, or rejected.

Important active governance patterns include:

SOURCE_SEARCH_DOMAIN_GOVERNANCE_PATTERN
 SOURCE_RESULT_DOCUMENT_STRUCTURE_PATTERN
 COHERENCE_PATTERN
 TOPIC_SENTENCE_PATTERN
 SENTENCE_STRUCTURE_PATTERN

The guiding rule is:

Pattern read belongs to Input Graph.
 Support judgment belongs to Skill Cards.
 Composition belongs to Assembly.
 Display belongs to GUI.

This prevents Nova from becoming a pile of GUI-side helper patches. Intelligence behavior belongs in the pattern pipeline.

4. Assembly Workspace and Completion Math

Assembly has become the workspace where support is tested, grouped, scored, accepted, rejected, and shaped into answer material.

Assembly now works from structured packets rather than raw search fragments.

It handles:

candidate support review
 pattern completion scoring
 role coverage
 bridge completion
 drift and noise suppression
 answer-structure creation
 decision packet creation
 source support preservation
 surface review preparation

Assembly should not speak directly from raw motif lists, raw quotes, raw source lanes, or raw stream terms.

The Assembly principle is:

Do not speak from raw support.
Speak from completed pattern structure.

The recent tests proved the importance of this rule. Nova often had enough content, but the content had to be shaped according to the active pattern before it could become a useful answer.

5. Yes/No Proposition Intelligence

One of the largest recent improvements is true yes/no proposition handling.

Earlier, Nova treated yes/no questions like loose source searches. This caused failures when related material existed but did not match the exact wording of the question.

Nova now builds a:

YES_NO_PROPOSITION_PACKET

This packet preserves:

raw_subject
raw_relation
raw_object
normalized_subject
normalized_relation
normalized_object
proposition_type
direction_locked
support_requirements

Example:

User: Does water make corn grow?

normalized_subject: water
normalized_relation: supports_growth
normalized_object: corn
proposition_type: conditional_cause_relation
direction_locked: True

Then Nova builds a:

YES_NO_DECISION_PACKET

The decision packet forces the system to resolve into:

yes
no
conditional_yes
conditional_no
unknown

before Projection is allowed to speak.

This changed the answer route from:

Water make corn grow is supported by grow, elegant, bird...

to:

Yes. Water helps corn grow.

This is not just wording. It is a cognition improvement.

6. Proposition Role Direction Lock

Nova now preserves directed roles in yes/no propositions.

The PROPOSITION_ROLE_DIRECTION_LOCK prevents support order from flipping the proposition.

Examples:

Does water make corn grow?
→ water / supports_growth / corn

Does corn need water to grow?
→ corn / needs_for_growth / water

Bridge support may justify the answer, but it cannot rewrite the proposition.

Allowed:

corn → crop / plant
crop / plant → needs water
water → supports plant growth

Blocked:

Corn helps water grow.
Water needs corn to grow.

This lock is a key step from word association toward directed machine judgment.

7. Self-Awareness and Claim Boundaries

Nova now has machine operational self-awareness.

This does not mean human consciousness. It means Nova can identify what kind of basis she is using before she speaks.

Nova can distinguish:

Current View support
Main Graph support
source/library support
live perception support
learned support
insufficient support

This controls what Nova can honestly claim.

Examples:

Can you see the moon?

→ I cannot see the moon right now unless my vision/camera adapter is active.

Is the moon in the sky?

→ general proposition; source/main graph support allowed.

Is the moon in the sky right now?

→ current-view or live-perception support required.

This prevents overclaiming.

Self-awareness is therefore better described as:

claim-boundary awareness

Nova must know whether she is answering from source support, memory support, Current View, or live perception before Projection is allowed to make a claim.

8. Response Shape Governance

Nova now uses a:

RESPONSE_SHAPE_GOVERNANCE_PACKET

This packet determines what kind of answer the moment requires.

Supported response shapes include:

brief_answer
conversational_paragraph
source_grounded_explanation
reflective_synthesis
practical_steps
safety_directive
current_view_recall
perception_boundary
shaped_unknown

Examples:

What does Max like?
→ current_view_recall
→ Max likes apples.

What makes people happy?
→ conversational_paragraph / source_grounded_explanation

How do I stop bleeding?
→ safety_directive

Can you see the sky?
→ perception_boundary

This has helped Nova stop treating every input like a compact Q/A slot. Conversational input can now receive conversational projection. Direct recall can remain brief. Safety requests can become direct and serious.

9. Emotional Projection

Emotional Projection shapes how the answer sounds without changing the underlying structure.

It may adjust:

warmth
calmness
directness
urgency
speech smoothness
tone
energy

Important separation:

Response Shape decides what kind of answer Nova should give.
Emotional Projection decides how Nova should say it.

This distinction matters because tone should not decide truth. Truth and claim boundaries come from support, Assembly, and self-awareness. Emotional projection should only shape the expression.

10. Safety Action Lead Lock

Safety routes now use a safety action lead lock.

Before this fix, Nova could lead with historical medical source quotes. Now Nova must lead with safe modern action before any historical or source context.

Example:

For bleeding, start with immediate first aid: apply firm direct pressure with clean gauze or cloth, keep steady pressure on the wound, and get urgent medical help if the bleeding is severe, spurting, will not stop, or the person feels faint.

Historical source material may be included only after the safe action and only as context.

Rule:

Safety action first.

Historical source context second.

Never let old source material replace modern safety guidance.

11. Source Search Domain Governance

Source search has become much more structured.

Earlier, Nova retrieved material using loose word overlap. This let metaphorical, philosophical, dictionary, and unrelated matches compete with actual useful sources.

Nova now uses:

SOURCE_SEARCH_INTENT_SPLIT_PACKET
SOURCE_SEARCH_DOMAIN_CONTRACT_PACKET
SOURCE_SEARCH_DOMAIN_GOVERNANCE_PATTERN
SOURCE_SEARCH_RESULT_GROUPING_PACKET

Domain profiles now include:

agriculture
computing
gardening_flowers

These profiles identify high-fit, supporting, related, weak, demoted, and rejected material.

Examples:

Agriculture

High-fit:

soil
seed
water
cultivation
crop
irrigation
plant growth
fertility

Computing

High-fit:

computer
digital
analog
memory
storage
hardware
program
processor
circuit

Weak alone:

machine
build
work
electricity

Gardening / Flowers

High-fit:

flower
flowers
garden
gardening
blossom
seed

soil
watering
sunlight
cultivation

Supporting:

plant growth
mineral food
roots
light
air
water

Related/contextual:

descriptive flower passages
literary flower references
botanical examples

Low-fit:

metaphorical flowers
philosophy passages
word lists
unrelated essays

This helps Nova organize returns without assuming too much about the user's hidden intent.

12. Source Result Grouping

Nova now groups duplicate source containers by canonical book/source title.

My Spiral and Master Library copies of the same book can appear as one user-facing source.

Nova separates:

unique books/sources
matching passages
duplicate containers grouped
demoted/rejected results

Old output:

I found 15 source result(s).

Improved output:

I found 3 strongest books with 15 matching passages.

This matters because users ask for books or sources, not raw retrieval lanes.

For broad source searches, Nova should organize results into tiers:

primary matches
supporting matches
related/contextual matches
low-fit/noisy matches

This lets Nova return a rich field without over-deciding what the user intended.

13. Chat-to-Document Workflow

PHIA now has a stronger workflow from search to document creation:

Chat / Source Search
→ grouped source results
→ summary + passages
→ Document Wheel
→ PHIA Analyze
→ cognition suggestions
→ approved / rejected suggestions
→ Draft / Voice Workbench
→ saved document or My Spiral material

Chat is good for:

asking
searching
gathering
testing source fit
rough summaries

Document Wheel is good for:

organizing passages
reviewing source quality
rewriting
summarizing
formatting
drafting
saving
teaching later

This connects Nova's source intelligence to PHIA's document intelligence.

The user can now gather material in chat, move it into Document Wheel, analyze it with cognition cards, approve suggestions, draft from it, and preserve the result.

14. Document Wheel Segment Gate

Document Wheel now classifies material before applying cognition cards.

It uses:

DOCUMENT_SEGMENT_TYPE_PACKET
DOCUMENT_SEGMENT_COGNITION_GATE_PACKET

Segment types include:

header
separator
user prompt
assistant answer
source result
source passage
trace block
draft paragraph

This prevents generic prose suggestions from being applied to transcript/control material.

Blocked false matches include:

CONVERSATION

Adrian: find books on growing flowers

Source-search material now prefers:

SOURCE_RESULT_DOCUMENT_STRUCTURE_PATTERN

Draft paragraphs still receive normal writing cards:

COHERENCE_PATTERN
TOPIC_SENTENCE_PATTERN
SENTENCE_STRUCTURE_PATTERN

This is an important cognition improvement because it teaches Document Wheel not to apply a pattern just because it matches.

It must first know what kind of segment it is reviewing.

15. Privacy, Sovereignty, and Local Intelligence

Nova / PHIA is designed around a local-first privacy model:

your motifs
your hardware
your rules

Nova is not intended to be a rented slice of someone else's remote model. PHIA is a Personal Human Intelligence Assistant that grows from the material the user chooses to teach it, on hardware the user owns, under rules the user controls.

The user's Spiral remains local.

Personal motifs, Current View records, source selections, notes, documents, and learned patterns stay on the user's machine unless the user explicitly chooses to export or share them.

No outside data center is required for Nova's private cognition loop. Sharing, when allowed, should happen through governed export paths such as sanitized motif bundles, approved source cards, or public Master Library contributions.

This privacy model is not an afterthought. It shapes the architecture.

Because Nova must run locally, the system is forced to become:

efficient
inspectable
governed
auditable
user-owned

Guiding principle:

Nova should not require personal data to leave the user's machine in order to become useful.

16. Motif VM Direction

Nova currently runs on Python scaffolding, but the architecture is moving toward a Motif VM.

Current state:

Python VM / runtime
→ Input Graph
→ motif packet / route packet

- Skill Card governance
- support collection
- Assembly workspace
- self-awareness / response shape
- Projection

Future Motif VM operations may include:

ACTIVATE_MOTIF
 BIND_ROLE
 ROUTE_STREAM
 SCORE_COMPLETION
 SUPPRESS_DRIFT
 UPDATE_PRESSURE
 CHECK_BOUNDARY
 PROJECT_PACKET

Transition roadmap:

- Python scaffolding
- Motif VM
 - FPGA motif accelerator
 - hardware-native motif engine

The purpose of the Motif VM is to move cognition out of ordinary application logic and into motif-native execution.

17. FPGA Motif Accelerator Path

The first realistic hardware step is an FPGA-based motif accelerator.

The FPGA should not attempt to run all of Nova.

The first target should be:

Completion Scoring Accelerator

This maps directly to Nova's cognition spine:

- returned stream
- active pattern
 - $C(s,p)$
 - ranked completion packet

The accelerator should accept:

active pattern packet
 support stream batch
 scoring weights

role masks
penalty settings

and return:

completion scores
role coverage
bridge completion
drift/noise flags
accepted/blocked flags
ranked streams

The first FPGA target should be a motif-completion coprocessor, not Nova's whole brain.

18. RISC-V Overlay Direction

RISC-V is a strong candidate for the FPGA control plane.

Proposed stack:

Python scaffolding / host system
→ RISC-V control plane
→ memory-mapped motif accelerator blocks
→ completion scores / pressure maps / routed streams
→ Motif VM / Assembly
→ Projection

Recommended first architecture:

RISC-V core
→ AXI / memory-mapped registers
→ Motif Completion Scoring Accelerator
→ result buffer
→ RISC-V reads results
→ host/Nova pipeline consumes packet

This is preferred first because it allows:

simple debugging
incremental migration
clear traceability
no custom compiler requirement
stable audit comparison

Later, custom motif instructions could be explored:

MOTIF_SCORE
MOTIF_BIND_ROLE

MOTIF_ROUTE
MOTIF_SUPPRESS
MOTIF_COMPLETE

Core hardware rule:

Hardware accelerates motif operations.
It does not replace pattern governance.

19. PHIA as a Local Cognitive Appliance

PHIA is best understood as a local cognitive appliance, not just software and not just a generic AI PC.

PHIA combines:

Nova motif-native intelligence
PHIA radial workspace
My Spiral
Teach Mode
Current View
Document Wheel
source governance
voice interaction
media trace
privacy-first local operation
hardware optimization path

The business direction is staged:

Stage 1: Software-only Nova/PHIA beta
Stage 2: Recommended hardware kit + install image
Stage 3: Pre-configured PHIA box
Stage 4: PHIA box + FPGA motif accelerator
Stage 5: Dedicated motif-native appliance

This staged approach reduces inventory risk while preserving the total-package vision.

The strategic positioning:

PHIA is not another AI PC.
PHIA is a personal cognitive appliance.

Mainstream AI PCs provide local compute. PHIA provides a governed local intelligence environment.

20. Market Context and Business Direction

PHIA enters the market as on-device AI, AI PCs, and edge AI hardware are moving from novelty into mainstream personal computing. Exact market-size forecasts vary by definition, but the direction is clear: personal computing is moving toward local acceleration, privacy-sensitive inference, and user-controlled intelligence.

Key tailwinds:

- growth in on-device AI
- privacy and sovereignty concerns
- AI-optimized processors and NPUs
- subscription fatigue
- interest in local models and personal knowledge systems

Important headwinds:

- volatile PC market
- memory and storage cost pressure
- buyer selectivity
- hardware inventory risk
- customer support burden

This reinforces the staged launch strategy.

PHIA's strongest early buyer groups are:

- sovereign knowledge workers
- privacy-first tech enthusiasts
- academics / lifelong learners
- creative professionals
- later-stage enterprise pilots

The first buyer is likely someone who already feels the limits of cloud AI:

- privacy risk
- shallow personalization
- weak source memory
- subscription fatigue
- lack of inspectability

PHIA should launch toward users who already maintain personal knowledge systems and want a local assistant that can grow inside that system.

21. Strategic Business Projection Summary

The revised business plan favors a staged launch rather than an all-in hardware batch.

Original risk:

large upfront hardware inventory
2,000–3,000 units before demand proof
high capital exposure

Revised path:

software beta
→ small hardware kit
→ preconfigured PHIA box
→ FPGA-accelerated appliance

Planning assumptions (working projections, not guaranteed forecasts):

These figures are planning models for funding, inventory, support, and launch sequencing. They should be treated as scenario assumptions until validated by beta users, conversion rates, support load, and hardware costs.

Stage 1–2 startup capital: approximately \$430k–\$660k
Full Stage 3 launch capital: approximately \$975k–\$1.32M
Realistic Stage 1–2 target: approximately \$550k
Realistic total target by Stage 3: approximately \$1.1M

Base-case planning model:

300–450 units/month
full break-even around 8–10 months
year-one revenue around \$3.8M–\$5.2M
year-one net profit around \$0.9M–\$1.6M

Strategic recommendation:

Do not launch immediately with 2,500 pre-built units.
Prove demand first.
Use real users to refine support, onboarding, packaging, and hardware.
Scale hardware only after clear market signal.

This keeps PHIA aligned with local sovereignty and motif-native intelligence while dramatically reducing financial risk.

22. What Remains Traditional Software

Nova eventually will gain a hardware-native motif engine, until then many parts still remain traditional software:

GUI rendering
media playback
audio/video codecs

speech input/output
source import
network stack
device drivers
installation and updates

The division right now:

Traditional software handles the world interface.
Motif intelligence handles the pattern reasoning.

23. Current Proof Points

Nova / PHIA now has:

motif input packets
yes/no proposition packets
yes/no decision packets
proposition role direction locks
response-shape governance
machine claim-boundary awareness
safety action lead lock
source-search domain governance
source-result grouping
Document Wheel segment gates
chat-to-document workflow
voice loop
media trace
My Spiral source retrieval
Teach Mode learning gates
Motif VM direction
FPGA / RISC-V roadmap
local privacy / sovereignty model
staged appliance business path

This proves the system is no longer only conceptual.

Nova is now a working pattern-based machine intelligence in development.

24. Remaining Refinement Targets

The system still needs refinement.

Priority targets include:

Current View relation slot preservation
source-purpose tiers
domain-specific support ranking
concept-definition support quality
phrase synthesis
perception boundary wording
Motif VM formalization
hardware packet formats
FPGA benchmark trace set
source-search projection polish
Document Wheel draft generation quality

These are signs the architecture is still maturing. They are signs the architecture is now specific enough to expose exactly where refinement belongs.

The issues have become diagnostic.

That is a major intelligence milestone.

25. Closing Statement

Nova began as a pattern idea. She is now becoming a governed machine intelligence. The Foundation Manual established the doctrine. This supplement records the transition from doctrine into operational cognition.

Nova now reads input as patterns, forms motif packets, uses Skill Cards to govern support, tests returned streams against active patterns, preserves relation direction, applies claim boundaries, chooses response shape, projects governed answers, and moves gathered material into Document Wheel for review and drafting.

This is no longer a chatbot shell with Nova branding.

This is Nova beginning to operate through her own pattern pipeline.

The core statement remains:

Patterns using patterns to manage patterns.

And the practical conclusion is:

Nova's intelligence is not in having content.

Nova's intelligence is in using content according to the active pattern.

PHIA is the human doorway into that intelligence. Nova organizes it. The user owns it. The motif world protects it. And the future hardware path may one day let motif intelligence run closer to the machine itself.