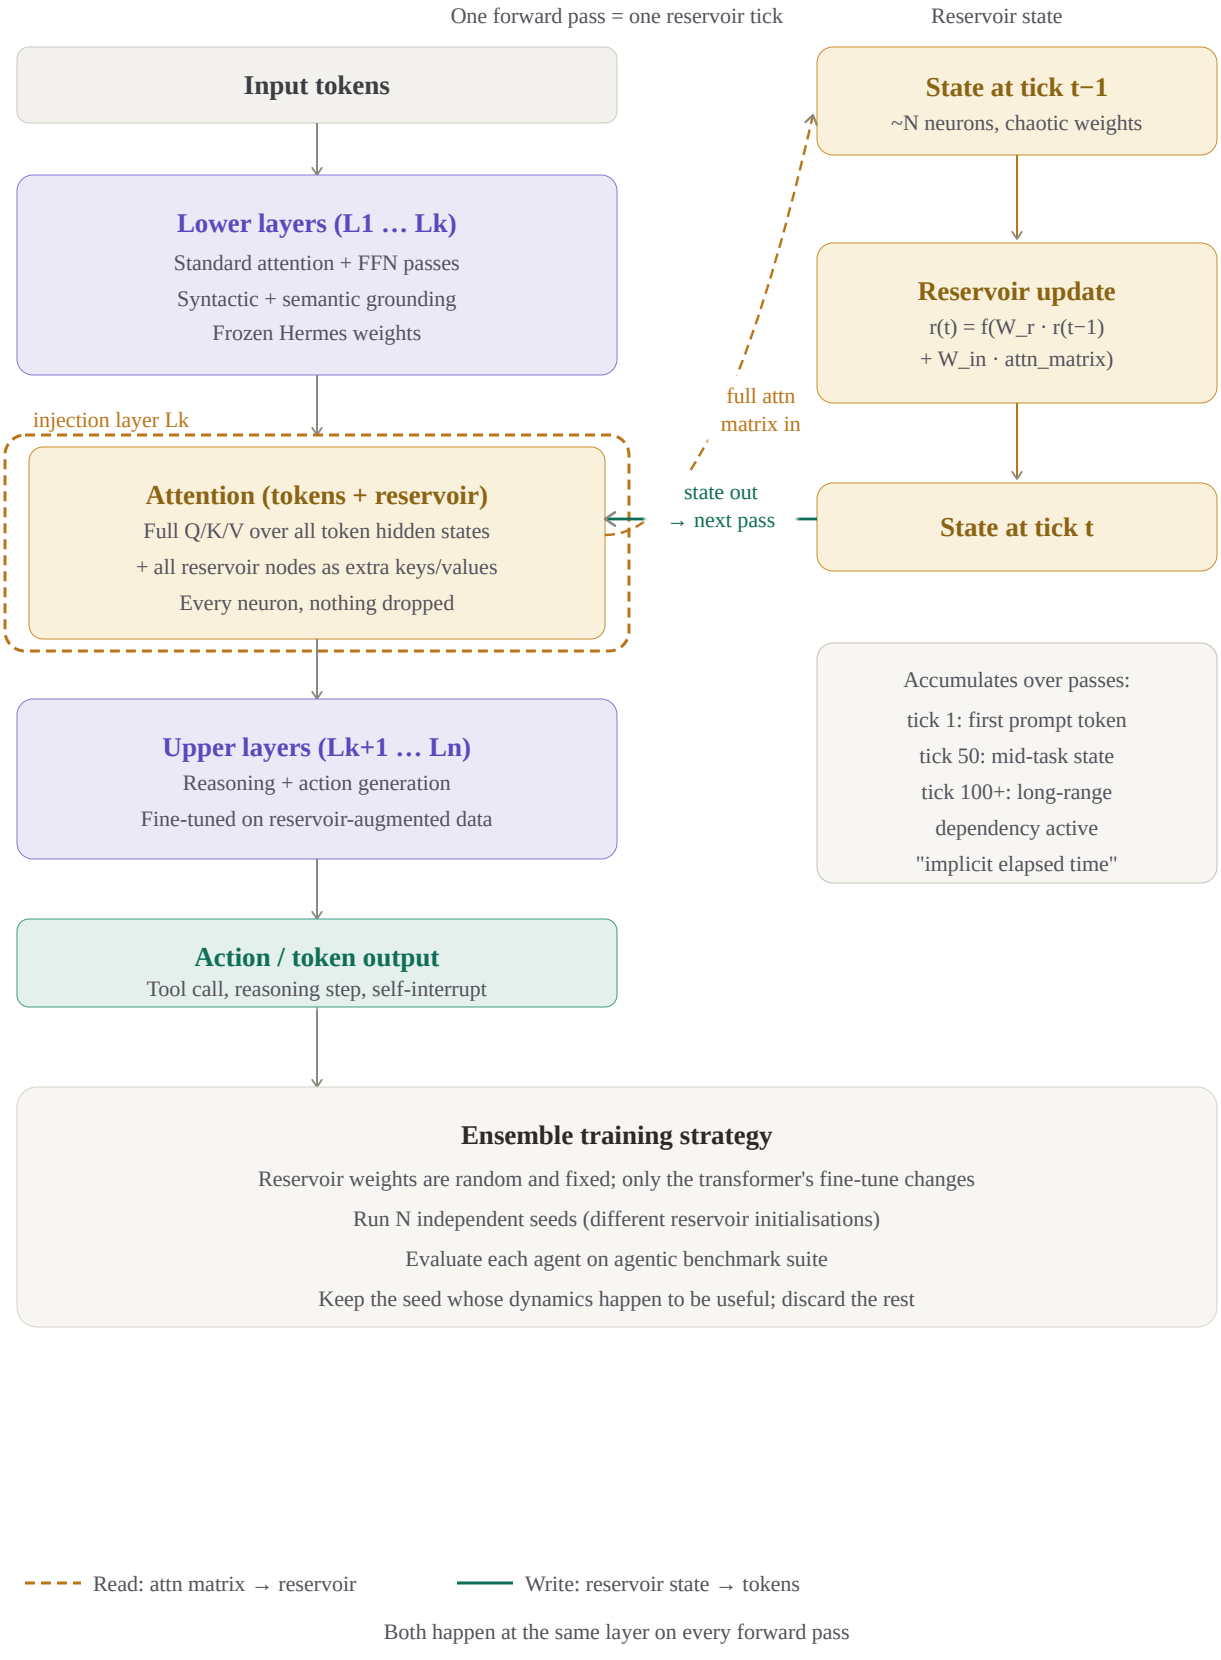

One forward pass = one reservoir tick
Reservoir state
Input tokens
Lower layers (L1 … Lk)
Standard attention + FFN passes
Syntactic + semantic grounding
Frozen Hermes weights
injection layer Lk
Attention (tokens + reservoir)
Full Q/K/V over all token hidden states
+ all reservoir nodes as extra keys/values
Every neuron, nothing dropped
Upper layers (Lk+1 … Ln)
Reasoning + action generation
Fine-tuned on reservoir-augmented data
Action / token output
Tool call, reasoning step, self-interrupt
State at tick t−1
~N neurons, chaotic weights
Reservoir update
r(t) = f(W_r · r(t−1)
+ W_in · attn_matrix)
State at tick t
full attn
matrix in
state out
→ next pass
Accumulates over passes:
tick 1: first prompt token
tick 50: mid-task state
tick 100+: long-range
dependency active
"implicit elapsed time"
Ensemble training strategy
Reservoir weights are random and fixed; only the transformer's fine-tune changes
Run N independent seeds (different reservoir initialisations)
Evaluate each agent on agentic benchmark suite
Keep the seed whose dynamics happen to be useful; discard the rest
Read: attn matrix → reservoir
Write: reservoir state → tokens
Both happen at the same layer on every forward pass