


SAFE AI NEEDS A CONTROL PLANE.




PTCP by Tensor Networks.

AI Is Becoming an Operator.



Past: Generating answers for human operators.




Present: Agents now call tools, access business data, write code, and initiate work across systems.

As capabilities grow, AI needs authority—but organizations need boundaries.

The Control Dilemma.

Prediction can propose a choice. It cannot grant extra authority.


To safely delegate responsibility, the agent's brain (prediction) must be entirely separated from the infrastructure's rules (enforcement).




The Paradigm Shift: From Implied Trust to Bounded Authority.

	Ungoverned AI		PTCP Governed AI	
Authority	Implied & Open		Bounded & Explicit	
Enforcement	Agent-controlled		Independent Network Controls	
Failures	Cascading & Widespread		Contained Locally	
Efficiency	Unpredictable Waste		Predicted Optimization	

The PTCP Operating Cycle.



Architectural Blueprint: Where PTCP Fits.




The Scale Principle: One Million Agents.

100 groups × 100 hosts × 100 agent identities.

Rule 1: Give each agent an identity
Attribute cost and access to specific tasks.

Rule 2: Keep decisions close to the work
Avoid central round-trips for ordinary actions.

Rule 3: Limit the scope of interruption
Contain incidents locally while preserving healthy work.



Safer Autonomy, More Useful Work.



CONTROL

Keep AI activity strictly within approved permissions and resource limits.



CONTAIN

Coordinate targeted intervention when activity exceeds limits, without total shutdowns.



OPTIMIZE

Use prediction to improve permitted decisions and pursue better useful-work efficiency.

Better Token Economics.

$(\text{Authorized Work}) \div (\$ \text{ Spend})$ = Better Token Economics.

More completed, authorized work per dollar. PTCP reduces wasted execution, avoidable waiting, runaway activity, and unnecessary interruption.



Make Control Part of Your AI Infrastructure.

