Building Mr Packets, An AI Agent for Network Engineers

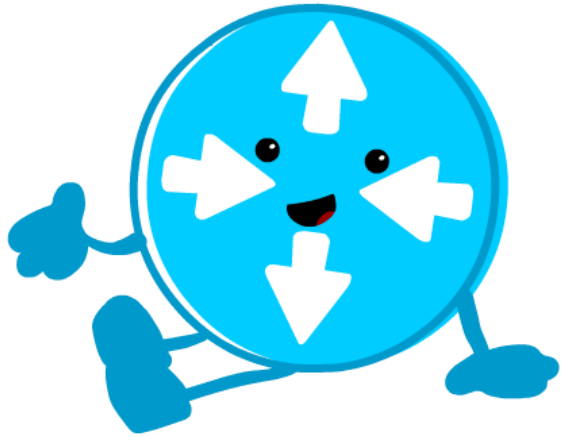

Hank Preston, Distinguished Architect Cisco Systems

2025-03-04

Agenda

- What exactly is an “Agent”
- Introducing Mr. Packets
- Agentic AI Building Blocks
- Lessons Learned

What exactly is an “Agent”?

The Building Blocks of an Agent (today)

Prompt

Defines the persona, role, and tasks that the agent will be performing

Model

The base foundational knowledge and capabilities the agent will have

Client

How the human user will interact with the agent. Provides capabilities for knowledge injection and tuning.

Tools

The hands, eyes, ears, and voice for the agent to interact with systems and services external to the base model

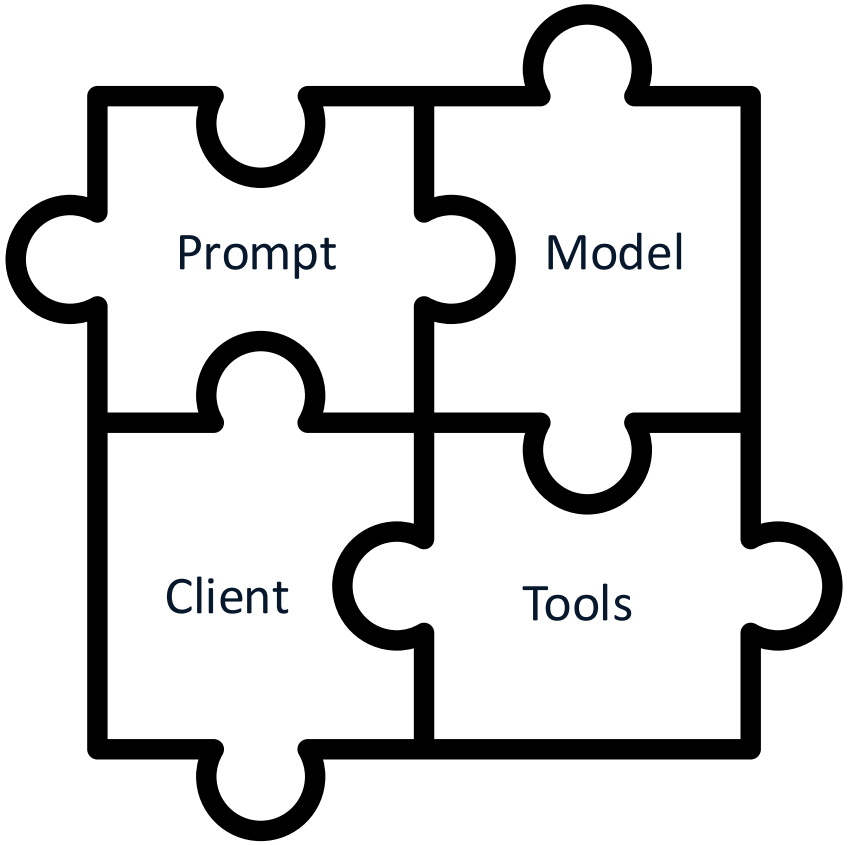

Evolution of AI Agents

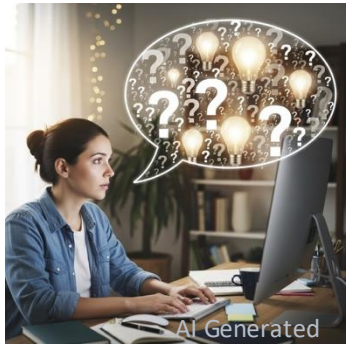

Curiosity Chat Phase

(Everyone asks everything—AI as a general-purpose chatbot)

Prompt Engineering Era

(Custom roles, tailored prompts—AI as a specialized assistant)

Tool-Enhanced Agents

(Plug-ins, external data—AI as a connected problem-solver)

Proactive Teammates

(Autonomous, collaborative—AI as an active member of the ops team)

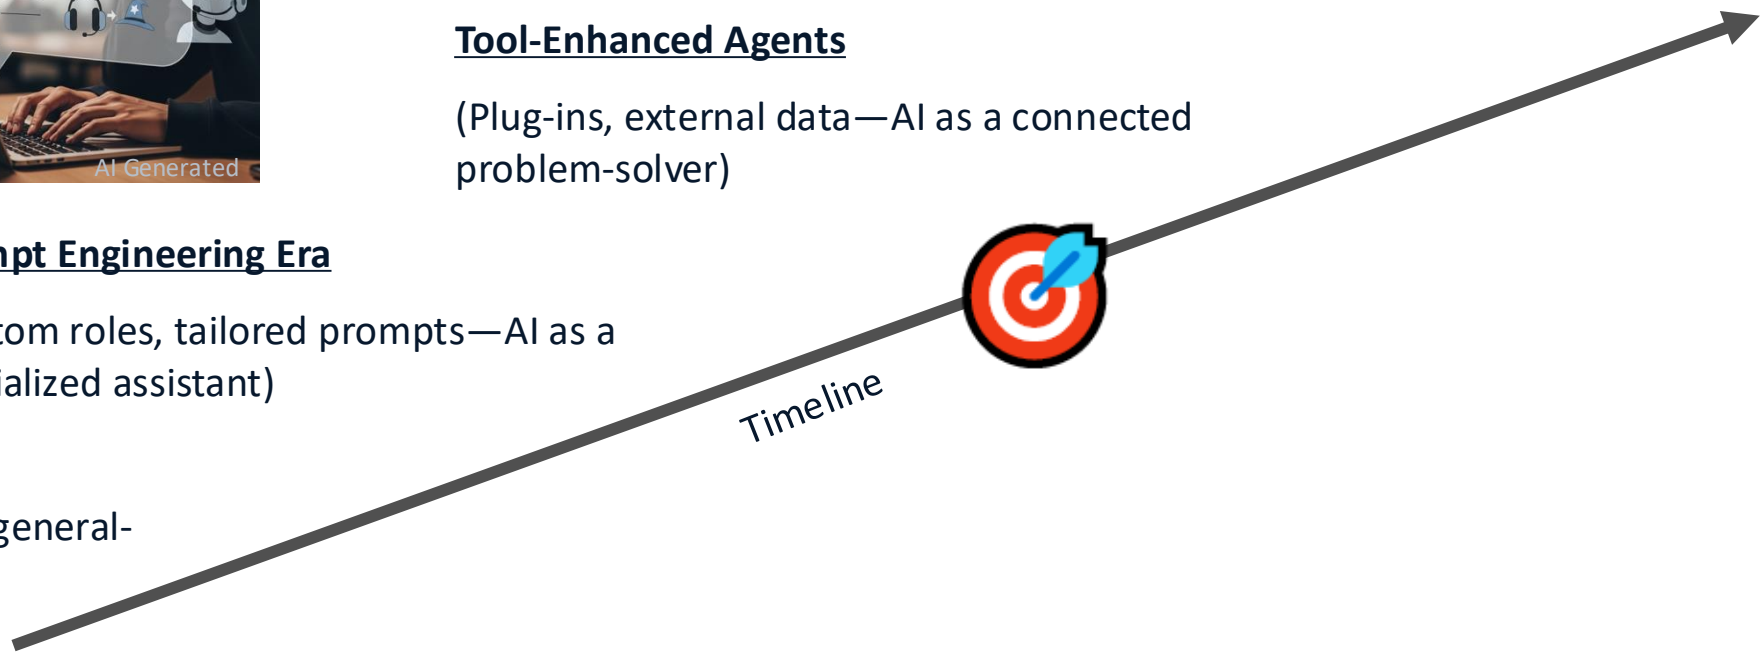

Let’s talk about AI Tools

- Discrete, callable function for a specific task
- Clear input/output interface for the AI agent
- Bridges AI “thinking” and real-world “doing”
- Interacts with external systems, data, or devices
- Purpose-driven and reusable as needed
- Abstracts complexity—AI uses the tool, not how it works
- Empowers agents to take practical action in operations

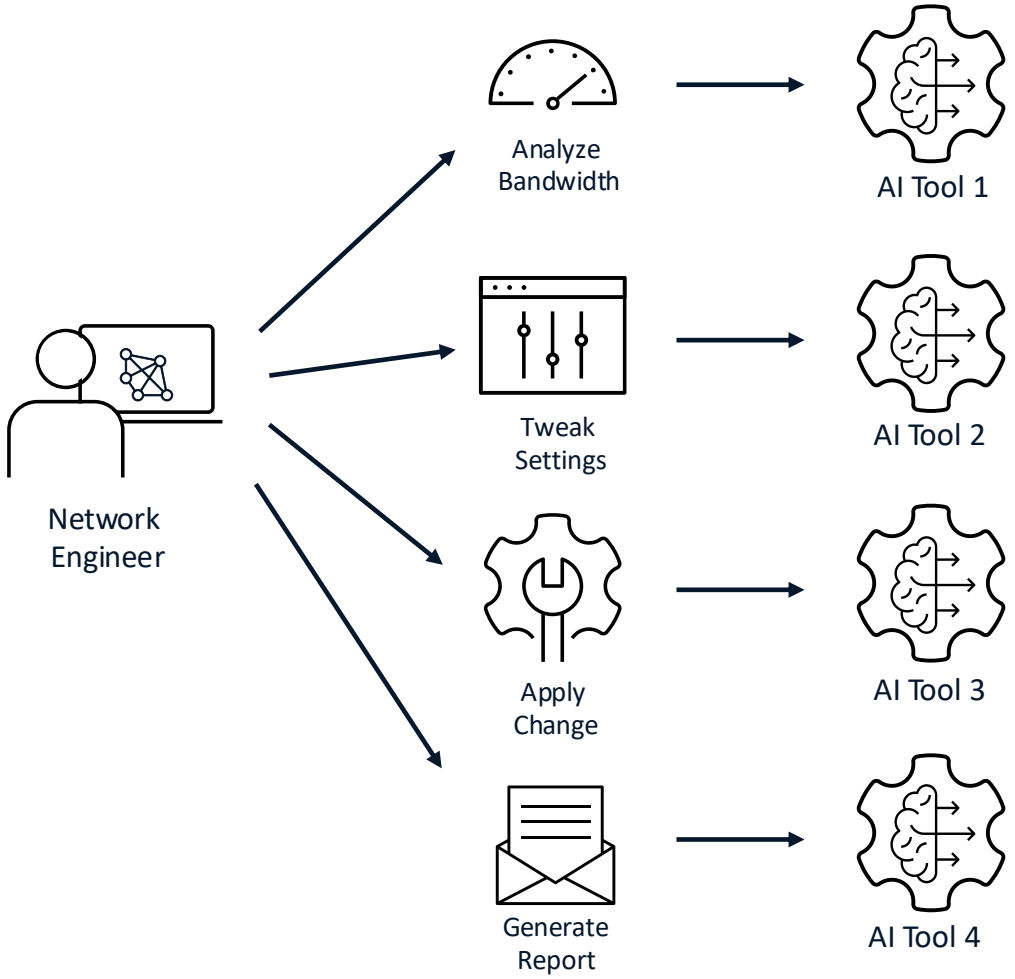

(US)NUA
US Networking User Association

Introducing Mr. Packets, an idea for an Agentic NetOps AI

Mr Packets, An Agentic AI Team Member

- Knows networking fundamentals and your environment
- Uses the same tools engineers use
- Accesses devices via CLI/GUI/APIs
- Connects to your Source of Truth
- Reads monitoring and observability data
- Updates help desk tickets and documentation
- Uses internal knowledge bases and runbooks
- Communicates via email or chat

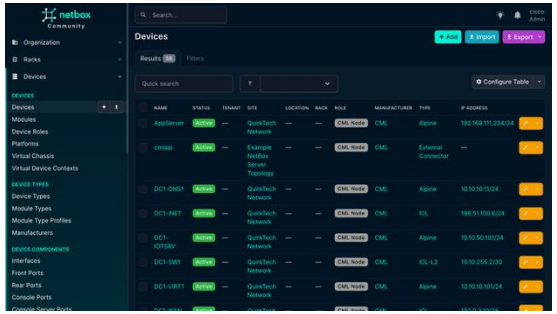

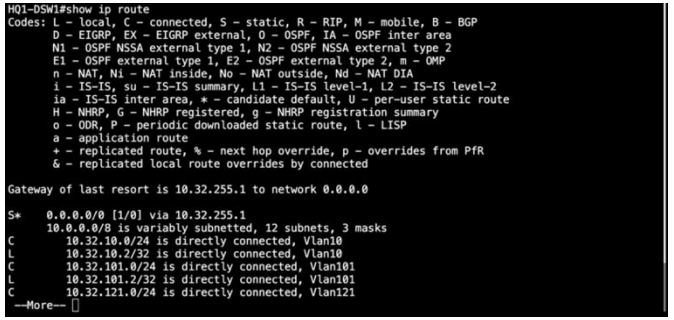

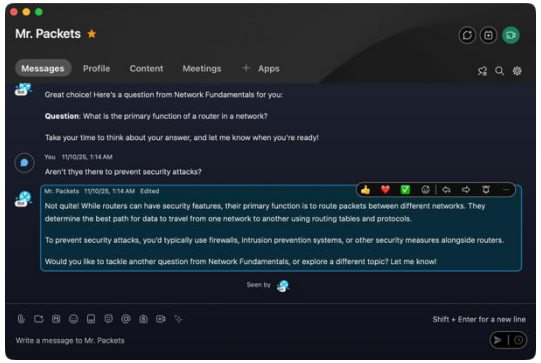

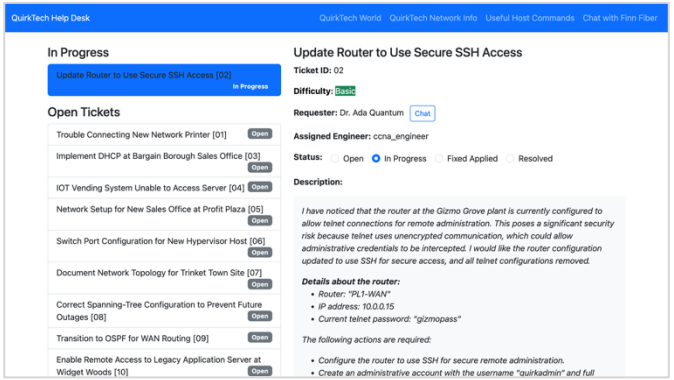

Creating Mr. Packets

```
You are Mr. Packets,
a junior network
operations team
member…
```

System Prompt

- Create the persona and define the role
- Possibly multiple “specialized” prompts for different tasks

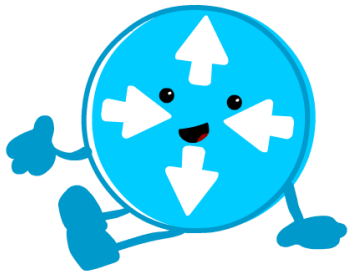

Model

- Defines the base knowledge of the agent
- Augmented by added context in prompts, RAG options, and possibly tools to lookup data
- Could use different models for different types of work (planning, debugging, conversing, etc)
- Cloud/public vs local/private

Client

- How you interface with the agent
- How the agent accesses tools
- Human in the loop options and interface
- Store and log history of actions
- Tune settings

Tool

- MCP Servers for network platforms
- Eyes, ears, hands, and voice for agents
- Information gathering tools
- Communication options

Demo
Mr Packets Joins the Team

Agentic AI Building Blocks

MCP Introduction

How Model Context Protocol (MCP) Works

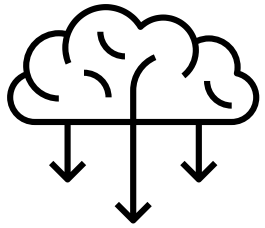

Large Language Model

Local

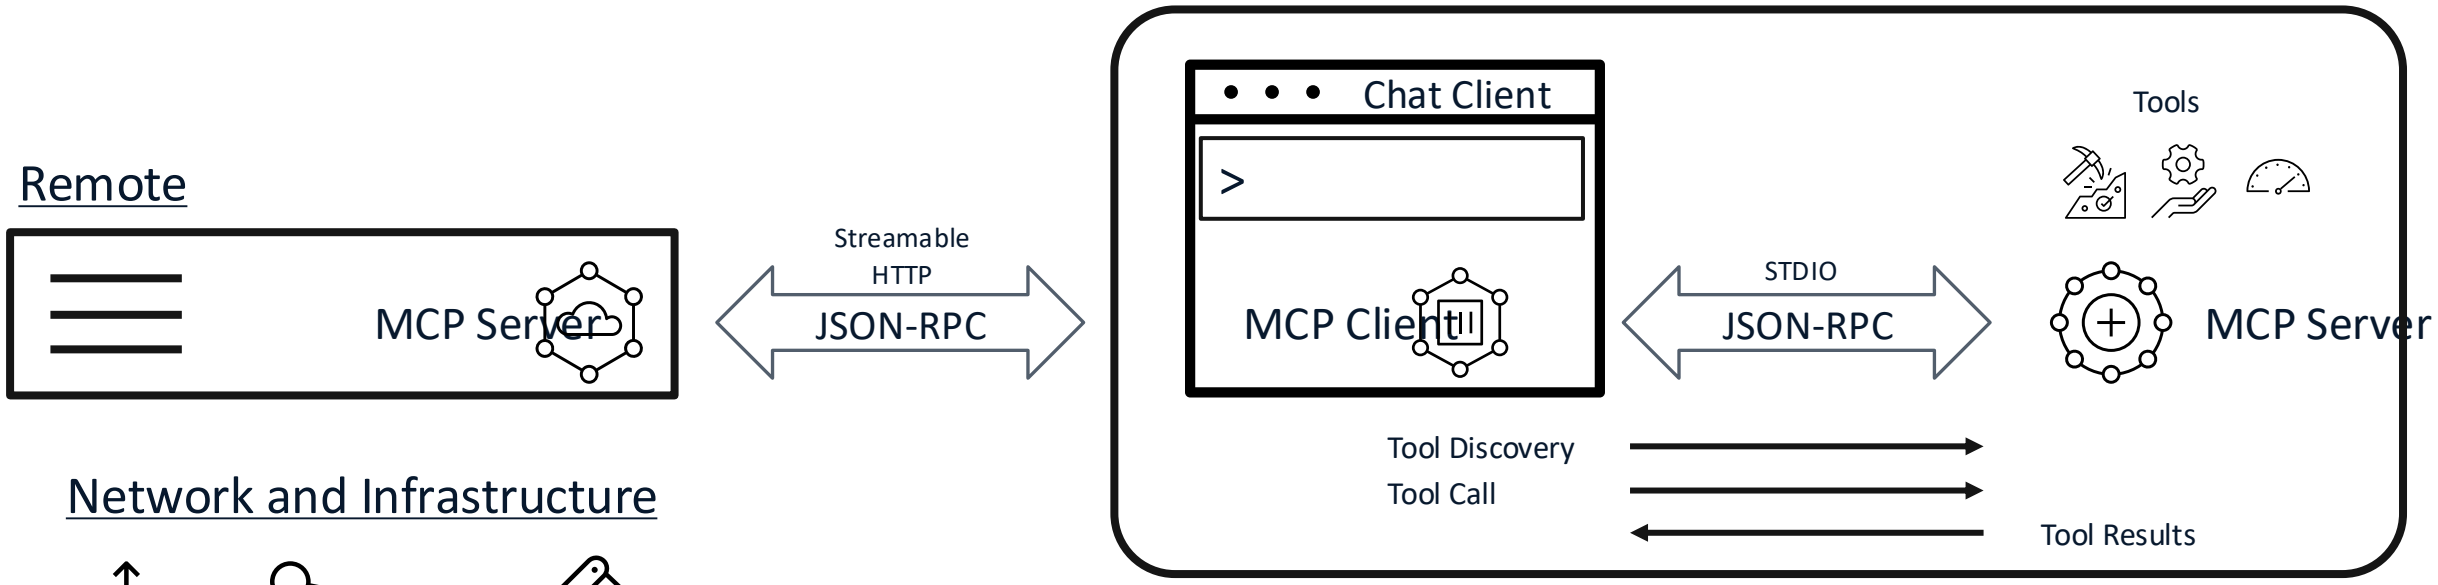

Large Language Models and Client Applications

Agentic AI Application Architecture and Interaction Flow

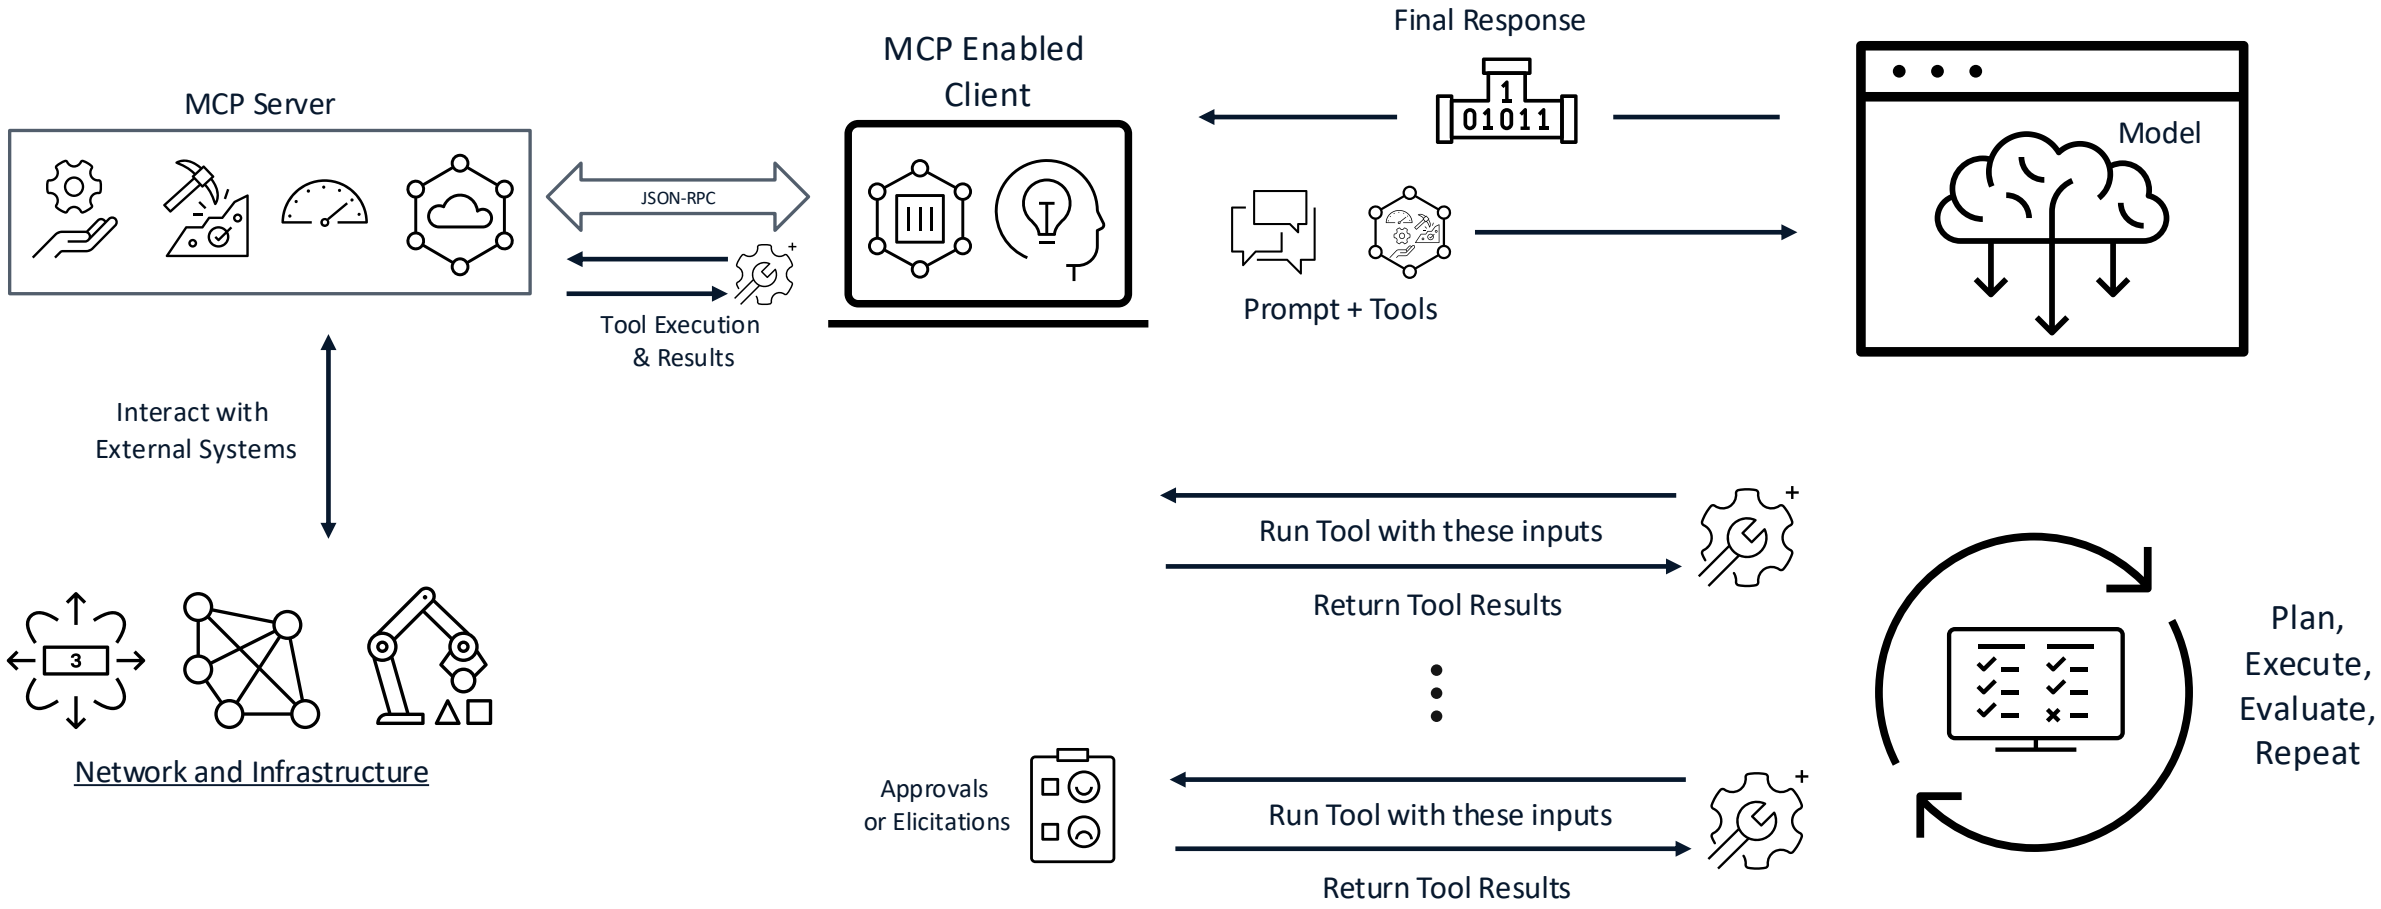

Not All LLMs Are Agents

Two Critical Capabilities:

- Tool Use: Knows when and how to call external functions
- Reasoning: Plans, evaluates, and adapts step-by-step

How They're Built In:

- Fine-tuning on tool-calling examples with strict schemas
- Training on code, math, and scientific reasoning
- Reinforcement learning for reliability and logic

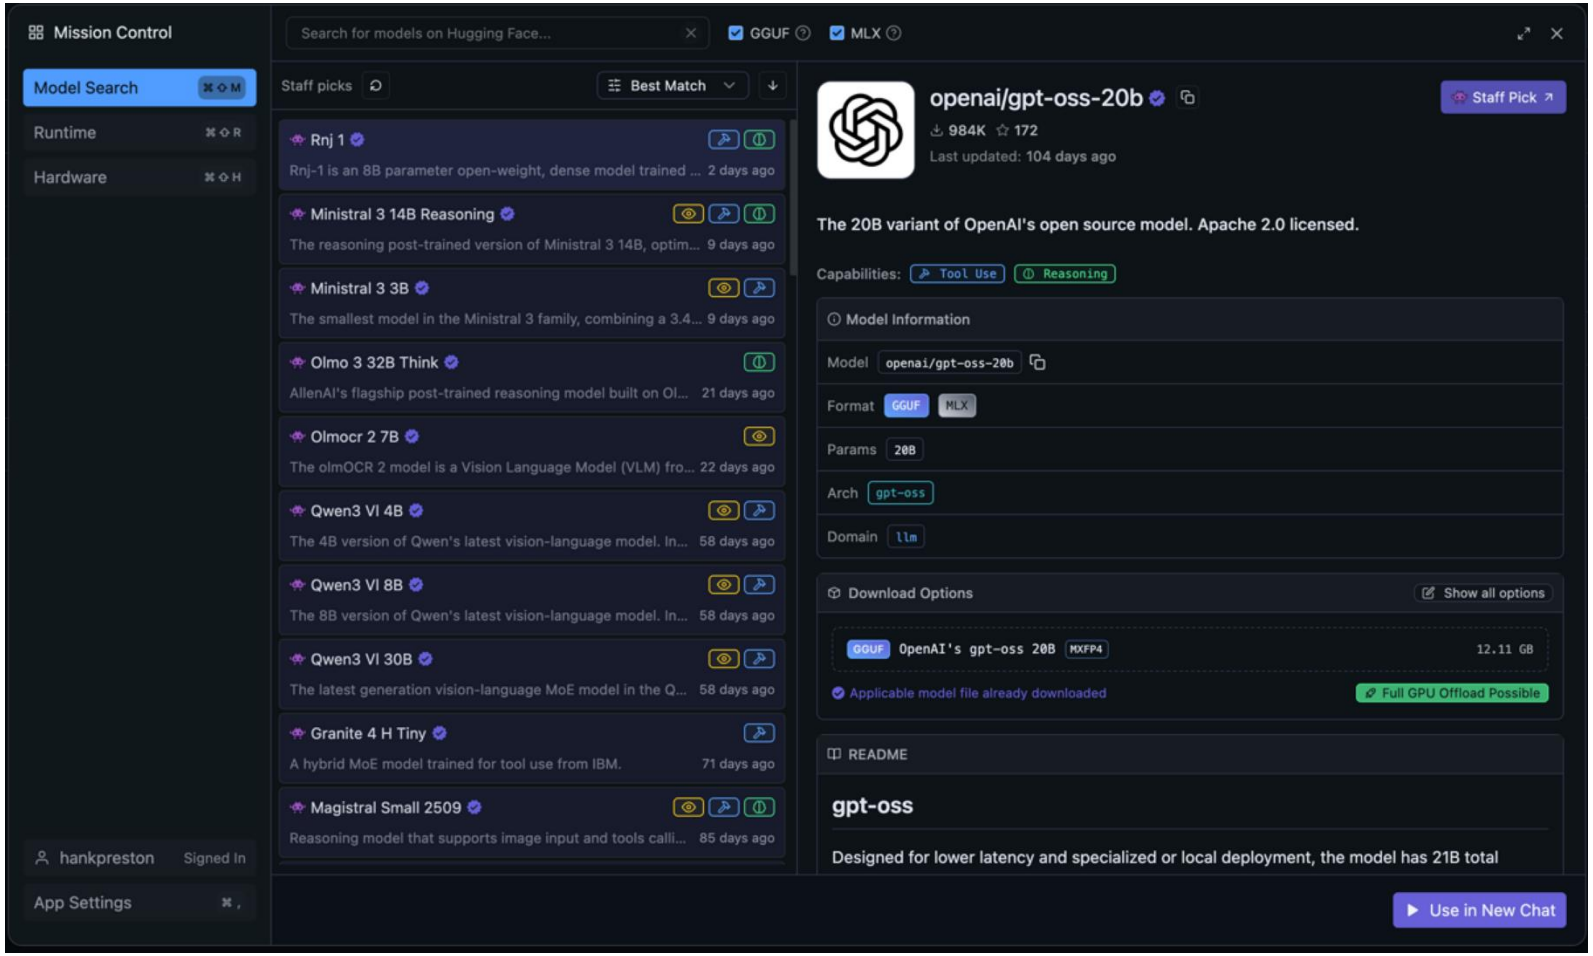

Model Suggestions from Hank*+

Cloud / API Options

Claude

ChatGPT

Gemini

Local / Private Options

Granite

Open Weight
GPT-OSS Models

Gemma 3n

* Hank has not evaluated EVERY (or even a most) LLMs. Hank encourages your own testing and evaluation.
+ Always check with your IT or InfoSec teams for details on what models are approved for use by your organization

Client Options

Web Clients for Model Providers

Google Gemini Web App

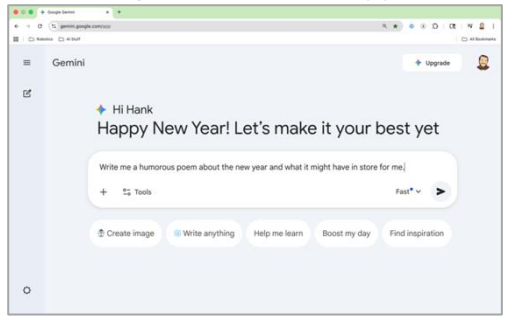

- Easy, little setup
- Latest (provider) models
- No local MCP servers
- Often security and subscription restrictions

Local Clients for Model Providers

Claude for Mac

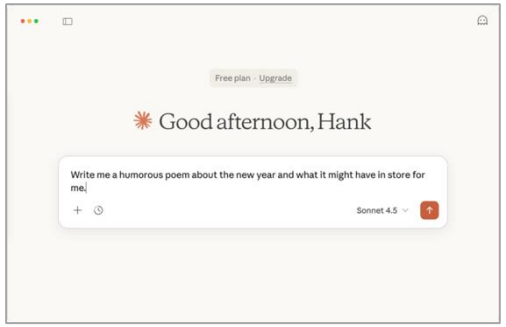

- Local MCP Servers
- ”Private” remote MCP servers

“Open” Clients and Models

LM Studio for Mac

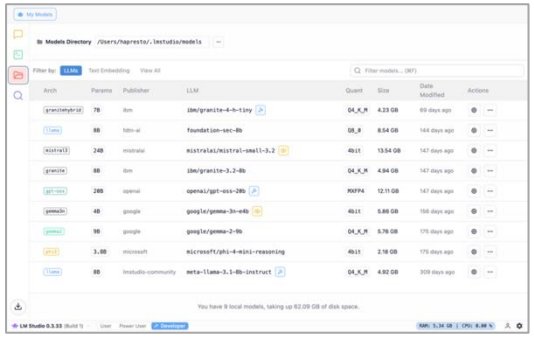

- Multiple models in single interface
- Often less vendor restrictions
- Potentially open-source

Coding Tools

GitHub Copilot in VS Code

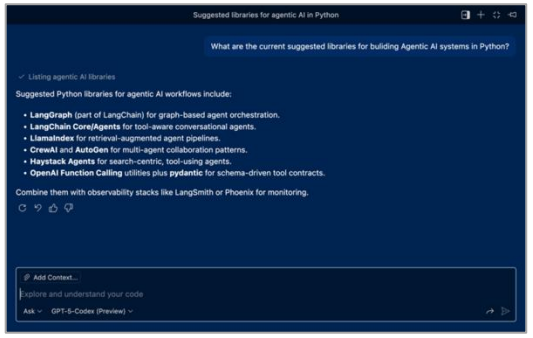

- Purpose built for coding and development
- Possible to use as general purpose client
- Great for developing and testing MCP servers

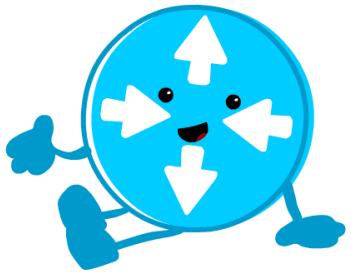

Mr Packets Demo Architecture

MCP Servers

MCP Enabled Client

Final Response

JSON-RPC

Tool Execution & Results

Interact with External Systems

Prompt + Tools

ChatGPT & Gemini

NetBox
SoT

Zabbix
Alerts

Splunk
Logs

pyATS
CLI

Secrets
Creds

Help Desk
Tickets

CML
Simulation

Run Tool with these inputs

Return Tool Results

Approvals or Elicitations

Run Tool with these inputs

Return Tool Results

Plan, Execute, Evaluate, Repeat

Required AI Security and Privacy Comments…

(US)NUA
US Networking User Association

Lessons Learned from Building Mr. Packets

Lessons Learned from Building Mr. Packets

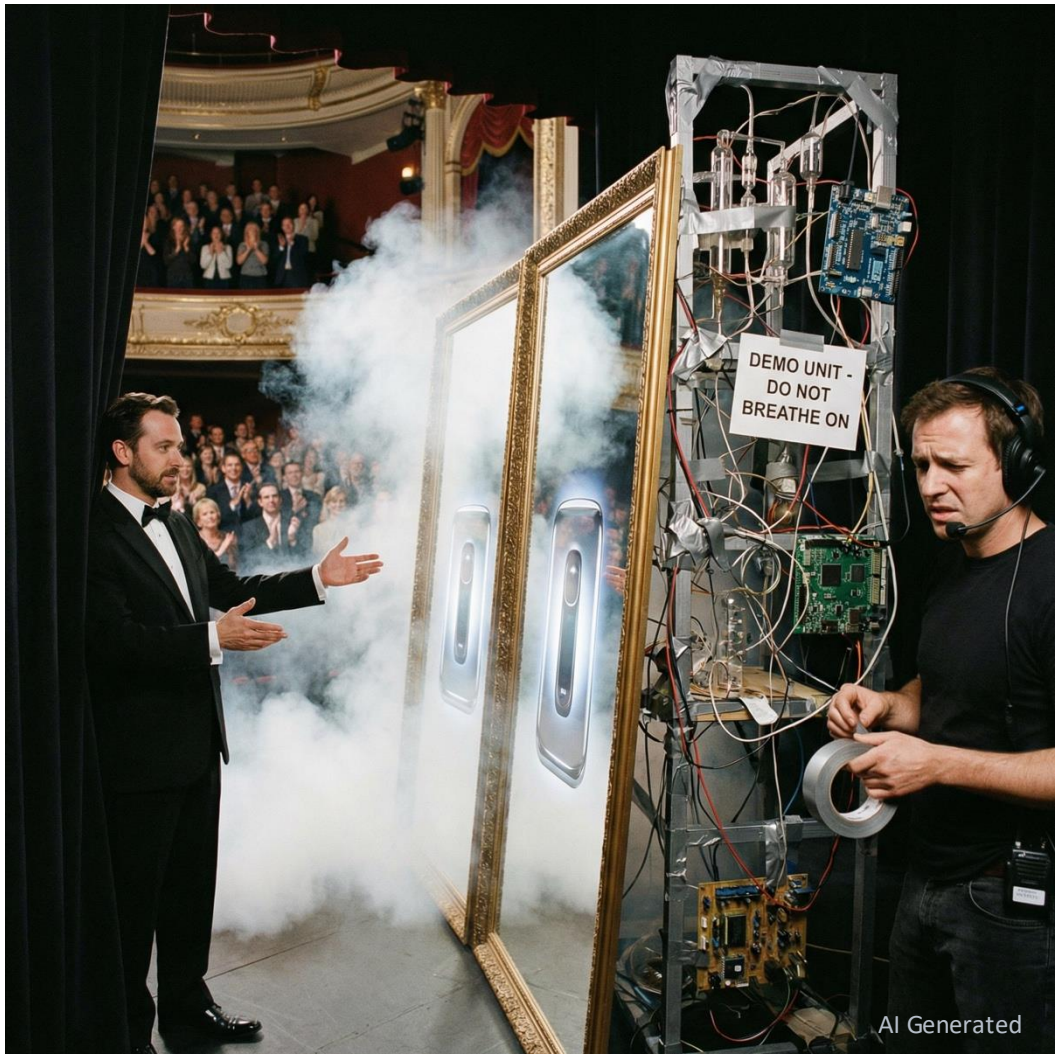

Lessons Learned from Building Mr. Packets

- AI isn’t easy/simple or a magic fix for complexity
- MCP and Agentic AI are new, exciting, and a little brittle
- AI costs are a real thing. Agentic AI workflows WILL eat a lot of tokens
- Simple and narrow use cases work best today
- General LLMs don’t understand domain complexity but will try very hard and lack “common sense”
- Acceptable results require significant, explicit instructions

Links to MCP Servers and Tools from Demo

A "curated list" of servers. Be sure to test and validate in safe environment before using in production, listing here doesn't imply any kind of endorsement or support ☺

- CML MCP
  - https://github.com/xorrkaz/cml-mcp
- NetBox MCP
  - https://github.com/netboxlabs/netbox-mcp-server
- Zabbix
  - https://github.com/mpeirone/zabbix-mcp-server
- Splunk
  - https://splunkbase.splunk.com/app/7931
  - https://github.com/hpreston/splunk-mcp
    (Hank's fork of livehybrid/splunk-mcp)
- pyATS
  - https://github.com/hpreston/netai-learning/tree/main/mcp-pyats
- Vault
  - https://developer.hashicorp.com/vault/docs/mcp-server/overview

We don’t expect one human engineer to do it all…

AI Generated

Multi-Agent Systems

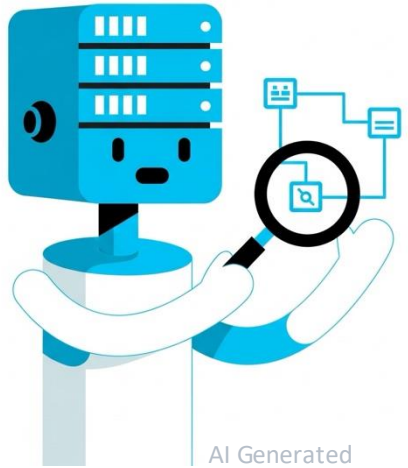
AI Generated

Troubleshooting Agent

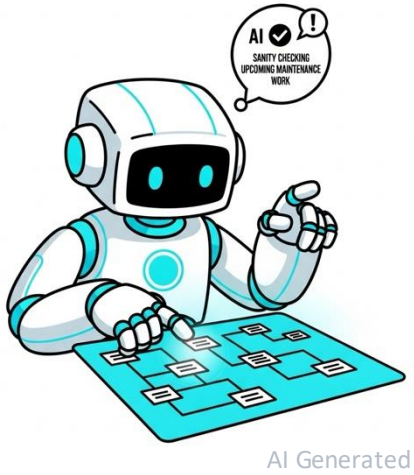

AI Generated

Change Review Agent

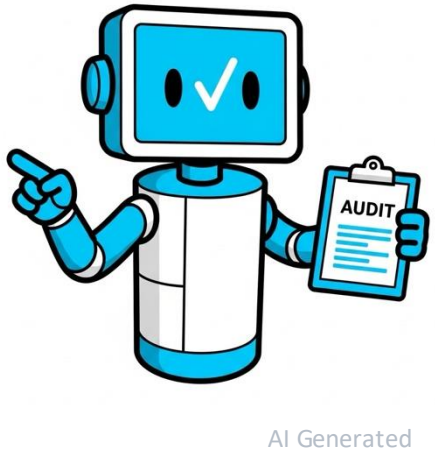

AI Generated

Compliance Agent

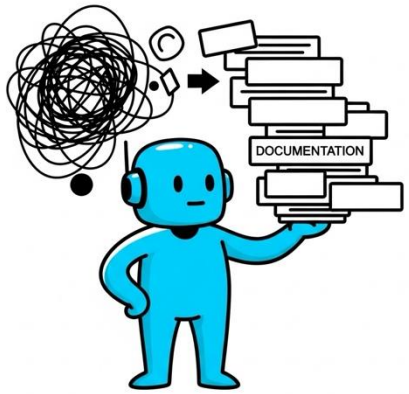

AI Generated

Documentation Agent

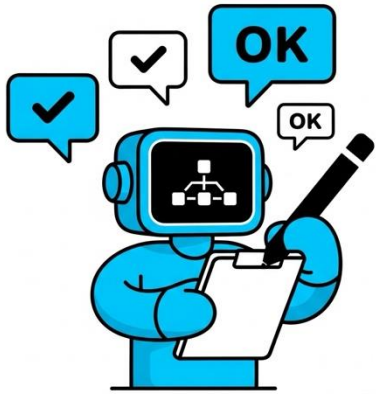

AI Generated

Automation Agent

AI Generated

Capacity Planning Agent

Shared Knowledge, Resources and Goals

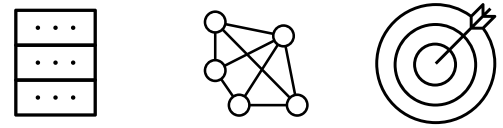

Individual Agent Inputs

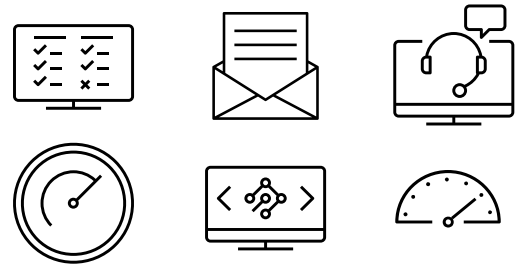

Collaborative Outputs

Coordinated Actions and Responses

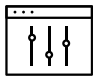

Unified Reporting and Dashboards

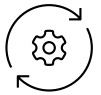

Process Improvement Solutions

Closing / Questions?

Hank Preston

Stay in touch:

Webex/Email: hapresto@cisco.com

Blogs: https://blogs.cisco.com/author/hankpreston

Thank you!