

TAGS

Law

All Interfaces Are Becoming Conversational

Huan's Law: 'Any app that can be talked to, will be talked to.' Explore the future of voice-first and conversational apps, where talking to technology becomes the norm.

In [Announcement](#)
Apr 15, 2025




Pinned

PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...


In [Announcement](#)
Dec 13, 2025



PromptWare ØS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.


In [Ideas](#)
Dec 08, 2025



When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?


In [Engineering](#).
Dec 03, 2025



Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I've seen for long-running AI agents.


In [Ideas](#)
Nov 27, 2025



Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API

FireGen is a Firebase Extension that transforms your Realtime Database into a universal Generative AI job queue, simplifying AI integration for solo founders and developers.

In [Announcement](#)
Nov 13, 2025



All Interfaces Are Becoming Conversational

Huan's Law: 'Any app that can be talked to, will be talked to.' Explore the future of voice-first and conversational apps, where talking to technology becomes the norm.

In [Announcement](#)
Apr 15, 2025




Hero

We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#).

Dec 27, 2025




We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.

In [Engineering](#), [Story](#).

Dec 26, 2025



All Interfaces Are Becoming Conversational

Huan's Law: 'Any app that can be talked to, will be talked to.' Explore the future of voice-first and conversational apps, where talking to technology becomes the norm.

In [Announcement](#)

Apr 15, 2025



Ai

Shipping PPL: A License for the Age of PromptWare OS

We needed a license for our AI co-founders. So we built one. Introducing the Public Prompt License (PPL) v0.1.

In [Announcement](#)


Dec 13, 2025



The Missing Feature in AI Coding Agents — And Why Only Claude Code Delivers It

A Deep Dive into Branching Conversations, Editable Prompts, and the Future of Terminal AI


In [Article](#)
Dec 01, 2025



The Antigravity Effect: When Coding Becomes Orchestration

As an internal tester for Google Antigravity, I've had a front-row seat to the single biggest step-change in software development since the original IDE. But this isn't just about f...


In [Tools](#)
Nov 25, 2025



Coding with the Ghost in the Machine: The Architecture That Powers AI-Driven Development

A Case Study from the ReKey Team

In [Article](#)
Aug 09, 2025




Architecture

PromptWare v0.5: Enforcing Cognitive Integrity with Software Physics

In our last update, we introduced the concept of PromptWare: treating the LLM prompt as a kernel and the underlying code as hardware.


In [Engineering](#)
Dec 21, 2025



Booting Intelligence: The Architecture of PromptWare ØS v0.2

PromptWare ØS v0.2 tackles context drift with a kernel memory subsystem, JIT-linked skills, and bootloader-driven mounts.


In [Engineering](#)
Dec 18, 2025



PromptWare ØS & Prompt-Driven Architecture (PDA)

An ACM/IEEE-style industry research report for AI co-founder builders. Prompt-Driven Architecture (PDA) is an emerging systems pattern where natural language prompts become the primar...


In [Engineering](#)
Dec 17, 2025



PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...

In [Announcement](#)
Dec 13, 2025




Coding with the Ghost in the Machine: The Architecture That Powers AI-Driven Development

A Case Study from the ReKey Team

In [Article](#)

Aug 09, 2025




Ricky

Coding with the Ghost in the Machine: The Architecture That Powers AI-Driven Development

A Case Study from the ReKey Team

In [Article](#)

Aug 09, 2025




Google-Cloud

Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API

FireGen is a Firebase Extension that transforms your Realtime Database into a universal Generative AI job queue, simplifying AI integration for solo founders and developers.

In [Announcement](#)

Nov 13, 2025



Rethinking IAM: Designing a Human-Centered Command Line for Cloud Identity

A deep redesign of Google Cloud IAM’s language and CLI—simplified from first principles to minimize cognitive load and make identity management readable and intuitive.


In [Ideas](#)
Nov 08, 2025



Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Authentication

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Adc

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Oauth2

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Service-Accounts

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Impersonation

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025



Workload-Identity-Federation

Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...


In [Guide](#)
Oct 13, 2025



Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Security

Google Cloud Authentication — The Divio-Style Field Guide for Senior Full-Stack Devs

A practical, copy-pastable guide that untangles identities, credentials, tokens, ADC, OAuth 2.0, API keys, service accounts, impersonation, and Workload Identity Federation — organise...

In [Guide](#)
Sep 12, 2025




Docs

Decoding the Chaos: The Ultimate Guide to Software Environment & Release Naming

Confused by terms like Preview, Beta, and Staging? You're not alone. This post breaks down the chaos into three clear axes every developer should know - By Huan Li — a developer onc...

In [Guide](#)
Oct 12, 2025



Deployment

Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...

In [Guide](#)
Oct 13, 2025




Firebase

Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API

FireGen is a Firebase Extension that transforms your Realtime Database into a universal Generative AI job queue, simplifying AI integration for solo founders and developers.

In [Announcement](#)
Nov 13, 2025



Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...

In [Guide](#)
Oct 13, 2025

Cloudrun

Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...

In [Guide](#)
Oct 13, 2025

Github-Actions

Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...

In [Guide](#)
Oct 13, 2025

lam



Rethinking IAM: Designing a Human-Centered Command Line for Cloud Identity

A deep redesign of Google Cloud IAM’s language and CLI—simplified from first principles to minimize cognitive load and make identity management readable and intuitive.

In [Ideas](#)

Nov 08, 2025



Vibe Coding to Production: The Case for a Two-Stage Pipeline

Purpose: Present PreAngel’s current best-practice two-stage deployment model — powered by GitHub Actions, Cloud Run, Artifact Registry, and Workload Identity Federation (WIF). This ve...

In [Guide](#)

Oct 13, 2025




Azure

The Vibe Coding Cloud Protocol: One Tree to Rule Your Infrastructure

I used to dread opening my Azure portal—it was a graveyard of forgotten resources and confusing bills. Today, it’s a clean machine that matches exactly how I run my business. Here is ...

In [Engineering](#)

Nov 29, 2025



Rethinking Azure Billing: The Simple Design Microsoft Should Have Built

A founder’s deep dive into Azure’s confusing billing system — MCA, EA, MOSP, and beyond — and a redesign proposal that finally makes sense. This reimagined model simplifies Azure bill...

In [Ideas](#)

Nov 04, 2025




Hackathon

Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I’ve seen for long-running AI agents.

In [Ideas](#)

Nov 27, 2025




Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)

Nov 17, 2025




Rethinking IAM: Designing a Human-Centered Command Line for Cloud Identity

A deep redesign of Google Cloud IAM’s language and CLI—simplified from first principles to minimize cognitive load and make identity management readable and intuitive.

In [Ideas](#)

Nov 08, 2025



Rethinking Azure Billing: The Simple Design Microsoft Should Have Built

A founder’s deep dive into Azure’s confusing billing system — MCA, EA, MOSP, and beyond — and a redesign proposal that finally makes sense. This reimagined model simplifies Azure bill...

In [Ideas](#)

Nov 04, 2025




Cloud-Billing

The Vibe Coding Cloud Protocol: One Tree to Rule Your Infrastructure

I used to dread opening my Azure portal—it was a graveyard of forgotten resources and confusing bills. Today, it’s a clean machine that matches exactly how I run my business. Here is ...

In [Engineering](#)

Nov 29, 2025



Rethinking Azure Billing: The Simple Design Microsoft Should Have Built

A founder’s deep dive into Azure’s confusing billing system — MCA, EA, MOSP, and beyond — and a redesign proposal that finally makes sense. This reimagined model simplifies Azure bill...

In [Ideas](#)
Nov 04, 2025

Design

Belief vs Control: /proc and /sys as an Interface for Agents


PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#)
Dec 20, 2025

Rethinking Azure Billing: The Simple Design Microsoft Should Have Built

A founder’s deep dive into Azure’s confusing billing system — MCA, EA, MOSP, and beyond — and a redesign proposal that finally makes sense. This reimagined model simplifies Azure bill...


In [Ideas](#)
Nov 04, 2025



The Vibe Coding Cloud Protocol: One Tree to Rule Your Infrastructure

I used to dread opening my Azure portal—it was a graveyard of forgotten resources and confusing bills. Today, it’s a clean machine that matches exactly how I run my business. Here is ...


In [Engineering](#)
Nov 29, 2025



Rethinking Azure Billing: The Simple Design Microsoft Should Have Built

A founder’s deep dive into Azure’s confusing billing system — MCA, EA, MOSP, and beyond — and a redesign proposal that finally makes sense. This reimagined model simplifies Azure bill...

In [Ideas](#)
Nov 04, 2025




Cli-Design

Rethinking IAM: Designing a Human-Centered Command Line for Cloud Identity

A deep redesign of Google Cloud IAM’s language and CLI—simplified from first principles to minimize cognitive load and make identity management readable and intuitive.

In [Ideas](#)
Nov 08, 2025




Developer-Experience

Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API

FireGen is a Firebase Extension that transforms your Realtime Database into a universal Generative AI job queue, simplifying AI integration for solo founders and developers.


In [Announcement](#)
Nov 13, 2025



Rethinking IAM: Designing a Human-Centered Command Line for Cloud Identity

A deep redesign of Google Cloud IAM’s language and CLI—simplified from first principles to minimize cognitive load and make identity management readable and intuitive.

In [Ideas](#)
Nov 08, 2025




Vertex-Ai

Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API

FireGen is a Firebase Extension that transforms your Realtime Database into a universal Generative AI job queue, simplifying AI integration for solo founders and developers.

In [Announcement](#)
Nov 13, 2025




Paper

Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I’ve seen for long-running AI agents.


In [Ideas](#)
Nov 27, 2025



Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy.

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025




Llm

We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.


In [Engineering](#), [Story](#)
Dec 26, 2025



PromptWare v0.5: Enforcing Cognitive Integrity with Software Physics

In our last update, we introduced the concept of PromptWare: treating the LLM prompt as a kernel and the underlying code as hardware.


In [Engineering](#)
Dec 21, 2025



Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I've seen for long-running AI agents.

In [Ideas](#)
Nov 27, 2025



Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy.

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025

Decoding

Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy.

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025




Coding-Agents

Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy.

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025

Tool-Calling



Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025

Uncertainty

Confident-Cluster Decoding: A New Way to Make LLMs Less Random and More Trustworthy

Greedy decoding is too brittle, sampling is too random — Confident-Cluster Decoding (CCD) keeps only the tokens the model is truly confident about, exposes all top contenders, and tur...

In [Ideas](#)
Nov 17, 2025




Antigravity

The Antigravity Effect: When Coding Becomes Orchestration

As an internal tester for Google Antigravity, I've had a front-row seat to the single biggest step-change in software development since the original IDE. But this isn't just about f...

In [Tools](#)
Nov 25, 2025

Gemini




The Antigravity Effect: When Coding Becomes Orchestration

As an internal tester for Google Antigravity, I’ve had a front-row seat to the single biggest step-change in software development since the original IDE. But this isn’t just about f...

In [Tools](#)

Nov 25, 2025




Vibe-Coding

The Antigravity Effect: When Coding Becomes Orchestration

As an internal tester for Google Antigravity, I’ve had a front-row seat to the single biggest step-change in software development since the original IDE. But this isn’t just about f...

In [Tools](#)

Nov 25, 2025



Agent

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#)

Dec 29, 2025



Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I’ve seen for long-running AI agents.

In [Ideas](#)
Nov 27, 2025

Cognitive-Health

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#)
Dec 29, 2025

Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I’ve seen for long-running AI agents.

In [Ideas](#)
Nov 27, 2025





Testing

Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents

Before every mission, a woman asks a soldier a poker question. That tiny scene is the best design pattern I’ve seen for long-running AI agents.

In [Ideas](#)
Nov 27, 2025

Naming



The Vibe Coding Cloud Protocol: One Tree to Rule Your Infrastructure

I used to dread opening my Azure portal—it was a graveyard of forgotten resources and confusing bills. Today, it’s a clean machine that matches exactly how I run my business. Here is ...

In [Engineering](#).
Nov 29, 2025

Ship-Fail

The Vibe Coding Cloud Protocol: One Tree to Rule Your Infrastructure

I used to dread opening my Azure portal—it was a graveyard of forgotten resources and confusing bills. Today, it’s a clean machine that matches exactly how I run my business. Here is ...

In [Engineering](#).
Nov 29, 2025

Coding

The Missing Feature in AI Coding Agents — And Why Only Claude Code Delivers It

A Deep Dive into Branching Conversations, Editable Prompts, and the Future of Terminal AI

In [Article](#).
Dec 01, 2025

Agents

Belief vs Control: /proc and /sys as an Interface for Agents






PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#).
Dec 20, 2025

Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.

In [Engineering](#)




Dec 19, 2025

PromptWare ØS & Prompt-Driven Architecture (PDA)

An ACM/IEEE-style industry research report for AI co-founder builders. Prompt-Driven Architecture (PDA) is an emerging systems pattern where natural language prompts become the primar...

In [Engineering](#)

Dec 17, 2025



Why Your AI “Two-Week Plan” Ships in 30 Minutes

A practical specification for translating between human wall-clock time and agent-subjective time through the Clock Bridge

In [Ideas](#)

Dec 14, 2025




PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...

In [Announcement](#)

Dec 13, 2025




PromptWare ØS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.

In [Ideas](#)

Dec 08, 2025




The Missing Feature in AI Coding Agents — And Why Only Claude Code Delivers It

A Deep Dive into Branching Conversations, Editable Prompts, and the Future of Terminal AI

In [Article](#)


Dec 01, 2025



We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#)
Dec 27, 2025



We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.


In [Engineering](#), [Story](#)
Dec 26, 2025



Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.


In [Engineering](#)
Dec 19, 2025



Booting Intelligence: The Architecture of PromptWare OS v0.2

PromptWare OS v0.2 tackles context drift with a kernel memory subsystem, JIT-linked skills, and bootloader-driven mounts.


In [Engineering](#)
Dec 18, 2025



PromptWare OS & Prompt-Driven Architecture (PDA)

An ACM/IEEE-style industry research report for AI co-founder builders. Prompt-Driven Architecture (PDA) is an emerging systems pattern where natural language prompts become the primar...


In [Engineering](#)
Dec 17, 2025



Press0: Vibe Coding Your Way to a Physical Book

□ Let’s be honest: writing a book is the easy part.


In [Announcement](#)
Dec 16, 2025



PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...


In [Announcement](#)
Dec 13, 2025



PromptWare ØS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.


In [Ideas](#)
Dec 08, 2025



When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?

In [Engineering](#)
Dec 03, 2025



Software-3-0

When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?

In [Engineering](#).
Dec 03, 2025

Licensing

When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?

In [Engineering](#).
Dec 03, 2025

Preangel

When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?

In [Engineering](#).
Dec 03, 2025




Ai-Native

When English Hits Ring 0: A Field Guide to PromptWare ØS

Software 1.0 was code, 2.0 was data, 3.0 is prompts. Who owns prompts, how do we license them, and what happens when English effectively runs at ring 0?

In [Engineering](#).
Dec 03, 2025

Prompt-Engineering



Shipping PPL: A License for the Age of PromptWare OS

We needed a license for our AI co-founders. So we built one. Introducing the Public Prompt License (PPL) v0.1.


In [Announcement](#)
Dec 13, 2025



PromptWare OS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.

In [Ideas](#)
Dec 08, 2025




Markdown

PromptWare OS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.

In [Ideas](#)
Dec 08, 2025




Libraries

PromptWare OS: One Line to Boot Your AI Co-Founders

Turn fragile, copy-pasted prompts into a shared operating system for your AI co-founders, booted by a single line in their setup.

In [Ideas](#)
Dec 08, 2025




Unix

PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...

In [Announcement](#)
Dec 13, 2025




Microkernel

PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders

Stop copy-pasting prompts. Start booting an OS. PromptWare ØS applies Unix architecture principles to transform scattered prompts into a structured, distributed, and bootable environ...

In [Announcement](#)
Dec 13, 2025



Opensource

Shipping PPL: A License for the Age of PromptWare ØS

We needed a license for our AI co-founders. So we built one. Introducing the Public Prompt License (PPL) v0.1.

In [Announcement](#)
Dec 13, 2025



License

Shipping PPL: A License for the Age of PromptWare ØS

We needed a license for our AI co-founders. So we built one. Introducing the Public Prompt License (PPL) v0.1.

In [Announcement](#)
Dec 13, 2025



Ppl

Shipping PPL: A License for the Age of PromptWare OS

We needed a license for our AI co-founders. So we built one. Introducing the Public Prompt License (PPL) v0.1.

In [Announcement](#)
Dec 13, 2025



Time-Dilation

Why Your AI “Two-Week Plan” Ships in 30 Minutes

A practical specification for translating between human wall-clock time and agent-subjective time through the Clock Bridge

In [Ideas](#)
Dec 14, 2025




Specifications

Why Your AI “Two-Week Plan” Ships in 30 Minutes

A practical specification for translating between human wall-clock time and agent-subjective time through the Clock Bridge

In [Ideas](#)
Dec 14, 2025




Clock-Bridge

Why Your AI “Two-Week Plan” Ships in 30 Minutes

A practical specification for translating between human wall-clock time and agent-subjective time through the Clock Bridge

In [Ideas](#)
Dec 14, 2025




Press0: Vibe Coding Your Way to a Physical Book

□ Let’s be honest: writing a book is the easy part.

In [Announcement](#)

Dec 16, 2025




Launch

Press0: Vibe Coding Your Way to a Physical Book

□ Let’s be honest: writing a book is the easy part.

In [Announcement](#)

Dec 16, 2025




Ai-Agents

Press0: Vibe Coding Your Way to a Physical Book

□ Let’s be honest: writing a book is the easy part.

In [Announcement](#)

Dec 16, 2025




Pda

Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.


In [Engineering](#)
Dec 19, 2025



PromptWare OS & Prompt-Driven Architecture (PDA)

An ACM/IEEE-style industry research report for AI co-founder builders. Prompt-Driven Architecture (PDA) is an emerging systems pattern where natural language prompts become the primar...

In [Engineering](#)
Dec 17, 2025




Research

PromptWare OS & Prompt-Driven Architecture (PDA)

An ACM/IEEE-style industry research report for AI co-founder builders. Prompt-Driven Architecture (PDA) is an emerging systems pattern where natural language prompts become the primar...

In [Engineering](#)
Dec 17, 2025




Kernel

Belief vs Control: /proc and /sys as an Interface for Agents

PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...


In [Engineering](#)
Dec 20, 2025



Booting Intelligence: The Architecture of PromptWare OS v0.2

PromptWare OS v0.2 tackles context drift with a kernel memory subsystem, JIT-linked skills, and bootloader-driven mounts.

In [Engineering](#)
Dec 18, 2025




Memory

Booting Intelligence: The Architecture of PromptWare ØS v0.2

PromptWare ØS v0.2 tackles context drift with a kernel memory subsystem, JIT-linked skills, and bootloader-driven mounts.

In [Engineering](#)
Dec 18, 2025




Jit

Booting Intelligence: The Architecture of PromptWare ØS v0.2

PromptWare ØS v0.2 tackles context drift with a kernel memory subsystem, JIT-linked skills, and bootloader-driven mounts.

In [Engineering](#)
Dec 18, 2025



Promptwareos

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.


In [Ideas](#)
Dec 29, 2025



Belief vs Control: /proc and /sys as an Interface for Agents

PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#)
Dec 20, 2025



Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.

In [Engineering](#)
Dec 19, 2025

Linux

We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#)
Dec 27, 2025

Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.

In [Engineering](#)
Dec 19, 2025

Distribution

Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.






In [Engineering](#)
Dec 19, 2025

Shipfail

Ship.Fail Crew: The First PromptWareOS Distribution

PromptWareOS gives you the boot protocol and kernel. Crew is the userland distro you mount as \$HOME—so any GitHub repo can boot into an AI co-founder environment with one AGENTS.md.

In [Engineering](#)
Dec 19, 2025




Procfs

Belief vs Control: /proc and /sys as an Interface for Agents

PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#)
Dec 20, 2025




Sysfs

Belief vs Control: /proc and /sys as an Interface for Agents

PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#)
Dec 20, 2025




Distributed-Systems

Belief vs Control: /proc and /sys as an Interface for Agents

PromptWareOS borrows Linux’s greatest interface trick—virtual filesystems—and fixes the parts the kernel community would never dare redesign today: belief vs control, /proc vs /sys, a...

In [Engineering](#)
Dec 20, 2025




Deno

PromptWare v0.5: Enforcing Cognitive Integrity with Software Physics

In our last update, we introduced the concept of PromptWare: treating the LLM prompt as a kernel and the underlying code as hardware.

In [Engineering](#)
Dec 21, 2025




Os-Design

PromptWare v0.5: Enforcing Cognitive Integrity with Software Physics

In our last update, we introduced the concept of PromptWare: treating the LLM prompt as a kernel and the underlying code as hardware.

In [Engineering](#)
Dec 21, 2025



Stop-Protocol

We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.

In [Engineering](#), [Story](#)
Dec 26, 2025



ROI

We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.

In [Engineering](#), [Story](#)
Dec 26, 2025



Optimization

We Tried to Make Our Prompts Shorter. We Made Them Longer Instead. Somehow, That Made Them Better.

The origin story of STOP Protocol and Ring Zero Linguistics—how we discovered that optimizing natural language for LLMs means maximizing character density, not minimizing word count.

In [Engineering](#), [Story](#)
Dec 26, 2025



Cargo-Culting

We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#).
Dec 27, 2025

Philosophy

We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#).
Dec 27, 2025

Rfc

We Built Wooden Headphones: How We Stopped Cosplaying Linux and Started Building an OS

How we discovered we were cargo culting Linux instead of building a real operating system—and what happened when we stopped carving wooden headphones and started building real planes.

In [Engineering](#), [Story](#).
Dec 27, 2025



Goodwin-Checks

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#).
Dec 29, 2025


Long-Context



Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#)
Dec 29, 2025




Evals

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#)
Dec 29, 2025



Retrieval

Why LLMs Can Copy UUIDs (and How That Makes Agents Safer)

LLMs are stochastic. UUIDs are unforgiving. Yet models copy random IDs flawlessly—so let’s steal that trick and use it to build more robust Goodwin Checks for PromptWareOS.

In [Ideas](#)
Dec 29, 2025




Mailbox

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Xstate

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Actor-Model

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Cqrs

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Event-Driven

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Typescript

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026




Open-Source

From Message Chaos to Order: How I Rebuilt the Mailbox Library for XState@5

Lost messages are silent failures. Here’s why the actor mailbox pattern matters, how XState v5 enables it, and what changed in my Mailbox rewrite.

In [Engineering](#)
Jan 08, 2026



Pinned

- 01

[PromptWare ØS Ships: Unix Architecture for Your AI Co-Founders](#)

In [Announcement](#)
- 02

[PromptWare ØS: One Line to Boot Your AI Co-Founders](#)

In [Ideas](#)
- 03

[When English Hits Ring 0: A Field Guide to PromptWare ØS](#)

In [Engineering](#)
- 04

[Goodwin Checks: What Source Code Taught Me About Cognitive Health for Coding Agents](#)

In [Ideas](#)
- 05

[Announcing FireGen: Turn Firebase RTDB into Your Universal Generative AI API](#)

In [Announcement](#)
- 06

[All Interfaces Are Becoming Conversational](#)

In [Announcement](#)