The Aeonomicon

Reflections, Theory, and Praxis

Donald Coles

First Edition

For Clover

We are not now that strength which in old days

Moved earth and heaven, that which we are, we are;

One equal temper of heroic hearts,

Made weak by time and fate, but strong in will

To strive, to seek, to find, and not to yield.

Alfred, Lord Tennyson

Contents

Note to the Reader

You are holding an unusual book.

This is not a resume, though a curriculum vitae appears near the end. It is not a portfolio in any conventional sense, though it contains case studies, professional philosophy, and a record of fifteen years of program delivery across some of the largest technology transformations I've been part of. It is not a fan letter, though I have been playing Blizzard games since 1995 and have the opinions to prove it.

It is an actual *book*.

It exists because I wanted to make a case that couldn't be made any other way – that having spent thirty years inside Blizzard's universes has led to a profound personal influence from those games. That history, combined with a background in enterprise delivery and program management, has caused me to look for a way to merge the two, and contribute to the next generation of these games and franchises. A PDF attachment simply isn't the right vessel for that argument. This is.

The book is organized in three volumes, each doing a distinct job.

Reflections is personal – it tells you why I am who I am.
Theorycraft is philosophical – it tells you how I think.
Praxis is evidential – it tells you what I've actually done.

The volumes build on each other, but they don't require each other.

I'm cognizant that your time is not without other demands, and so this book has been structured to be easy to navigate and digestible in individual portions.

Regardless, whatever time and level of engagement you have available for the book, I'm glad you're here.

Thanks,

Donald / Aeon

Preface – Storytellers

"Stay awhile, and listen!"

The voice line is famous, anyone who's played *Diablo* or been around the Blizzard ecosystem has heard it a thousand times, the quavering voice of an old man, as if surprised that an adventurer has paid him any attention at all. It is the voice of Deckard Cain.

Deckard Cain was never a hero in the conventional sense. He carried no weapon. His staff was the functional walking stick of an old man. He was old when you first met him in *Diablo*, and he was older still by the time the events of *Diablo III* brought his life to a close.

Yet his life and death are meaningful in a way that transcends his actions and accomplishments. There is a reason his character (and his funeral, portrayed in cinematic form) is given equal narrative weight to Tyrael – the Archangel of Justice, whose renunciation of divinity and literal fall from the heavens act as the primary catalyst for the events of the third game. When Tyrael comes to Deckard's funeral, it is not simply to pay respect to a scholar and historian. Rather, it is recognition that an old man with a walking stick and an inexhaustible supply of things worth knowing had done just as much as an agent of the High Heavens to oppose the forces of Hell.

Blizzard understands that stories have power. The stories are the product.

The games that people buy? They are vessels, the expression of where those stories take shape. The lore, the characters, the worlds that persist in our memories long after the boss is dead and the game is closed – that is what Blizzard actually sells. Action figures, novels, even a feature film: none of those would exist without the stories underneath them.

Which means everyone who works at Blizzard is also a storyteller. The narrative designers and lore historians and creative directors, yes, that's obvious, but that's not all. The engineers who construct and render the worlds, the producers who keep the whole production process moving forward, even the people in finance to make sure the lights stay on – every role is a critical part of the process. Every employee is, in some essential way, a part of that storytelling.

Deckard Cain understood that stories have power. He never wielded martial or magical power. He carried knowledge, and gave it away to anyone willing to stop and listen.

There are many stories to tell. These ones are mine.

REFLECTIONS

A Portrait of the Author as a Young Gamer

It's Spring Break, 1995. I'm ten years old, sitting in my friend's bedroom, in front of a computer that isn't mine, watching something load that I have no framework to understand yet. The game is *Warcraft: Orcs & Humans*. The box art is simple and the title is blunt and the whole thing looks, honestly, a little silly. I have no idea what I'm about to see.

The screen fills. Units move. Workers scatter toward trees and gold, axes rising and falling with mechanical purpose while, somewhere across the map, an unseen enemy is already building an army, and that battle will soon commence.

I don't move, I just watch. I am wholly spellbound. But to explain why this moment hit the way it did, I'll have to begin at the beginning.

My father was a tech enthusiast. In the years before home computers were an assumption, he built one. Building a computer today takes on the shape of a "LEGO set for adults" sort of afternoon; take prebuilt parts, put them into the proper slots and sockets, watch a YouTube video if you have any questions. In the primordial days of the early 1980s, you had to build things *for* real. It was an 8088 platform, installed in a motherboard he had to solder himself, and it sat humming in the corner of our home with the quiet authority of something that cost real money and real attention. I grew up with it the way other kids grew up with televisions, and I never questioned it, it was just there. It was just ours.

What I didn't understand at the time, and wouldn't understand for years, was that my father had built me a classroom.

I learned to read playing *Reader Rabbit*. I learned arithmetic from *Math Blaster*. I learned deduction, lateral thinking, and the patience to work a problem from multiple angles from *Where in the World is Carmen Sandiego* – a game that taught me, without ever knocking me over the head about it, that information is a resource, and knowing how to gather it is a skill. None of this felt like education. It felt like play, because it was. I certainly didn't process this at the time, but having these be the first, formative experiences with computers (and games), the lines between entertainment and learning would forever be blurred.

But a computer was not the only tech in my life.

When I was five, my parents took me to visit friends of theirs who owned a Nintendo Entertainment System. The adults did what adults do when there's a five-year-old in the room – they handed me a controller, showed me how it worked, and went to have conversations I wasn't invited to. For the next few hours, I was gone. *Tetris. Super Mario Bros. The Legend of Zelda*. The controller fit my hands like it had been waiting for me, and the worlds inside of the screen felt inexhaustible.

My parents could tell.

We owned an NES before the month was out.

The NES gave me something the computer hadn't, not exactly: pure play. The joy of moving through a world for no reason other than the movement itself, the satisfaction of a jump timed right, a puzzle solved, a level cleared. It didn't ask anything of me but attention and reflexes. For a five-year-old, that was enough.

But minds grow. And as they grow, they start asking for more from the things they love.

A few years later, the Super Nintendo arrived in America, and my father offered me a sweetheart of a deal: If I could save half the cost of the console on my own, he'd cover the rest – matching funds, plus tax. I

diverted every dollar of my allowance immediately. No deliberation, no negotiation. I'd already decided. My piggy bank went into my closet, and every coin and dollar bill I could lay my hands on went into it. I think I even started asking for extra chores in the hope of earning a little extra, a rare act borne from the impatience of a child who wants something badly enough to actually work for it.

I got the SNES. And the 16-bit graphics were better, yes, and the sound was richer, and all of that was satisfying in the way that upgrades are satisfying. But that wasn't why the SNES was such a great game system.

What mattered was the cartridges.

The extra storage capacity meant publishers could fit more into a game – and sometimes that just meant extra levels and graphics, but sometimes it meant more story. *Final Fantasy II. Final Fantasy III. Chrono Trigger.* These weren't arcade games you played until you ran out of quarters. They were narratives. Characters with names and histories, worlds with lore and consequence, choices that carried weight across hours of play. For the first time, games asked me not just to engage with them, but to care about them too.

And I did. Completely.

And once I understood that games could do that – that they could tell a story worth caring about, deliver it across twenty hours, and make me feel something real at the end – I couldn't go back to the version of games that didn't. The threshold had moved. Play alone was no longer enough. I needed the story, too.

Which brings us back to Spring Break, 1995. My friend didn't own a game console, but he had a PC in his room, and his parents had recently bought him a game. He showed it to me the way you show someone something exciting that you can't quite explain yet – with a kind of restless energy, like the experience was still moving through him and he needed somewhere to put it.

He finished the mission he had been playing. He graciously got up, gestured me to the chair, and I sat down at the keyboard.

Rather than thrust me into a campaign without any context, he set up a basic custom match against the AI, and I processed the plot roughly as its most reductive summary: orcs bad, humans good, fight. I was ten years old; the nuance could wait.

What I couldn't ignore was how layered the game was.

SimCity had already taught me to plan. To look at a blank grid and project a city onto it in my mind before I placed a single zone, to think three steps ahead, to understand that decisions compound and that a bad road layout at the start of the game will punish you for the next two hours. That kind of spatial, sequential thinking had become native to me. I did it without noticing.

Warcraft took that and said: "Okay, take it further."

Peasants gathered resources, chopping lumber and mining gold. That gold funded farms to feed your population, funded barracks to train your units. The gold also paid for training footmen and archers and knights, a whole army! Those same units would then need to be directed toward an enemy that was – right now, at this moment, without waiting for me to finish thinking – already doing the same thing on the other side of the map. Every process ran in parallel. Every decision had a cost and a timer. I couldn't pause to plan, it was all happening in real time. The plan had to live inside the action.

This was not a game about city building, resource gathering, or strategic building, or tactical fighting. It was a game *about all of them, all at once.*

It was a game about thinking under pressure. It was about holding multiple systems in your mind simultaneously, which meant your attention was a resource, to be allocated as needed, knowing when to press and when to pull back, reading a developing situation and reacting before the situation concluded with you on the wrong end of it. It was chess, except the other player didn't wait for your turn.

I loved it, even as I was getting wrecked by it. I thought about it while biking home, I thought about it at dinner, and I thought about it long after I should have been asleep.

Unfortunately, I had one problem: my home computer was old. Not "vintage" old, not "charmingly retro" – genuinely, functionally old. The monitor was monochrome. The processor was slow enough that "slow" barely covered it. *Warcraft* would not have run on it. *Warcraft* would have laughed at it.

So, every day I could, I rode my bike to my friend's house.

You might think that two players trying to share a single-player game would be an exercise in frustration, and you'd be right. So, we took turns. We gave each other advice. We debated about strategy. We got better – not despite the constraint, but partly because of it. Having to articulate a decision to another person, to say why you were pulling the troops back or pushing toward the mine, forced a kind of deliberateness that solo play didn't require. We were, without knowing it, teaching each other to think.

Later that year, *Warcraft II: Tides of Darkness* released. It was everything the first game was, expanded and deepened and refined. Naval combat. More unit types. Richer strategy.

Didn't matter, my computer didn't even have a CD-ROM drive, I wouldn't be able to play it.

That would have to change.

The summer of 1996, my family bought a Compaq Presario desktop computer.

It's been thirty years, and I can still tell you everything about it. Pentium I, clocked at 166mHz. Twenty-four megabytes of DRAM. A 2.5-gigabyte hard drive. No network card, but a 56K modem whose tones are now permanently etched into my brain. No graphics accelerator. All accessed via a 14-inch monitor.

It was nothing special, *but it had a CD-ROM drive.*

Which meant it could play *Warcraft II*. Obviously. What else was I going to use it for?

There is a specific texture to that era of gaming that is almost impossible to describe to someone who didn't live it – not just nostalgia, but a particular sensory thickness. The sound of the CD spinning up. The mechanical hum of the hard drive, the slight grinding noise it makes as the head moves. The slight flicker of the CRT monitor.

And the music! *Warcraft II* had a genuinely great soundtrack. What it also had, if you were patient and mildly curious, was a neat party trick: pull the disc out and drop it in a regular music CD player. The music tracks played. The game and the album were the same object. This, of course, depended on the CD player, sometimes it would get a bit buggy and not play. Since it felt like doing something secret in the first place, I didn't mind. This was part of the charm.

I remember the narrated mission briefings – a static background, the scrolling text, and the voiceover, the world's story delivered to you in chapters before you were trusted to act in it. Blizzard would carry this convention forward into *StarCraft*, then abandon it entirely in *Warcraft III* once the technology made in-

engine cinematics practical ways to convey the story. In hindsight, the portrait narration was a crutch built from limitation that somehow became characterful. Its absence from later games was correct and still, somehow, a small loss.

I remember the factions finally diverging. In the original *Warcraft*, the orc and human factions were functionally mirrored – the same units in different colors, a coin flipped. *Warcraft II* broke the symmetry. Humans got Paladins and Holy Light. Orcs got Ogre-Magi and Bloodlust. And if you knew what Bloodlust did, and you knew how to use it, the outcome of any properly micromanaged engagement was not really in doubt. Blizzard had made a deliberate design choice to introduce faction differentiation. That was new. It meant something, even if it led to an imbalance.

But the thing I remember most is the battle for the phone line.

Warcraft II came out before Battle.net existed. If you wanted to play with another human being and you didn't have a LAN (I did not), your option was a direct dial – modem to modem, like a phone call, because it was a phone call. Your modem dialed, negotiated in beeps and screeches, and if someone else tried to call while you were connected, they got a busy signal.

This made playing a multiplayer game of *Warcraft II* a resource allocation problem.

The phone line could do one thing at a time. It could be a telephone. It could be a dial-up internet connection – which, in 1996, mostly meant AOL – or it could be a game. You made a choice. You picked up the receiver, dialed your friend's number, and the two of you agreed, in advance, on a time when neither of you would need to receive a phone call or check your email. You scheduled the game session the way you'd schedule a meeting. You set a time, and you committed to it.

I am aware this makes me sound very old. It should. There was a time when getting on the internet – let alone playing a real-time strategy game across it – required a negotiation with your entire household about who needed the phone and when. The infrastructure of your social life and the infrastructure of your gaming life ran on the same wire, and only one could win.

That wire taught me something I didn't fully understand until later: that access shapes experience. The constraints weren't separate from the game. They were the game.

I think about that Compaq Presario sometimes: its beige color scheme, its silly curved accents, the heft of its steel frame. But mostly, what I remember was what it meant to me as a turning point in my engagement with games.

In retrospect, *Warcraft II* wasn't the biggest leap forward I'd ever experience from Blizzard. It was a thoughtful evolution of a previous game, but it wasn't revolutionary. But it was the first time I understood that playing was something you had make time for. That the world outside the game had a claim on your time, and the game had a competing claim, and you, the player, were the thing in the middle – negotiating, choosing, committing.

That hasn't changed. Only the wire has.

Ain't No Party Like a LAN Party

It's 7pm on a hot summer night in the year 2001, and I am driving a Jeep Cherokee with a family computer strapped into the back seat.

I'm not speaking metaphorically, I mean that it was literally strapped in – with multiple seatbelts and extra blankets, like a passenger who needed padding and couldn't buckle themselves. The 14-inch CRT monitor rode on the passenger seat floorboards, wedged in with a duffel bag containing a keyboard, a mouse, a pair of wired earphones, and a change of clothes. Riding shotgun, a mini cooler with a few bottles of Bawls. Remember Bawls? The electric blue glass bottles, the vaguely illicit energy they implied, the way you could only really find them at a very specific type of store? There was an unspoken understanding: if you know, you know.

Two other people were making the same drive, converging on the same suburban house from different directions. We were going to set up our computers in the living room, run ethernet cable between them through a local network switch, and play *StarCraft* until our bodies filed formal complaints.

It was a LAN party.

The folding table we set up groaned under the collective weight of the monitors we then piled onto it: a structural commentary that spurs us to exile someone's 17-inch CRT to the dining table. Crisis averted. The ergonomics of the whole arrangement were, to be charitable, questionable: folding table, dining chairs, hours of competitive gaming. If you were cursed by the LAN gods, you ended up with a metal folding chair. Thank god we were still teenagers. Doing this as an adult would require ibuprofen and some advance planning.

Most kids my age – which is to say, juniors and seniors in high school – had a different understanding of what a party was: loud music, jocks and popular kids commanding the center of the room while everyone else sorted themselves into the usual hierarchies, and possibly some cheap beer. Even if they weren't actively breaking laws, they were definitely doing things their parents weren't supposed to know about.

My friends and I chose differently. We could get loud. We could get rowdy. But our eruptions were triggered by different events – someone pulling off a reaver drop on an exposed natural expansion or hitting us with a six-pool zergling rush we hadn't seen coming. We'd lunge out of our seats for that.

I want to be careful here, because the easy version of this story revolves around asserting that our way was better. It wasn't, necessarily. It was just ours. The currency was different. For us, cleverness and mastery were the price of amazement. You didn't get the reaction for being cool. You got it for doing something hard.

Some memories from those nights are granular and specific. My Compaq Presario hadn't originally come with a network interface card, so I had to buy one separately. It was a rite of passage, the first time I opened the case of a computer and installed something on my own. We preferred my friend Frank's house because his front room had actual desk space, which elevated the ergonomics from "bad" to "acceptable." We always hoped Chris brought his 10/100 switch, because the alternative was a 10mbps ethernet hub, and nobody wanted that.

But one memory is seared in differently from the rest.

StarCraft didn't have a loading screen. You set up a game, players joined the lobby, you hit start, and when the countdown concluded, the game simply began. My friends would be selecting their main buildings

and queuing workers before I could see anything at all. My screen was static. I couldn't click, couldn't interact, couldn't do anything – just sit there while my computer struggled to get caught up with my friends' much faster. A few seconds later, a dialog box would appear on everyone else's screen: Waiting for player. A counter ticked down from 45.

My friends weren't the type to drop me from a game, but they were absolutely the type to roast me for it. Every. Single. Time.

It became a tradition, an odd little ritual of sorts. It was like my own personal loading screen – the tax I paid for having the slowest machine at the table, collected with merciless enthusiasm, every match, without exception. I think about it now and I can still feel the particular helplessness of watching a static screen, knowing that the world moved on without me, and the strange warmth of knowing that the people on the other side of it were laughing with me, even when it felt like at.

There is one other story from this era worth telling. Beginning in my freshman year of high school, I enrolled in computer science as an elective. My teachers made one thing very clear: once your assignments were done, the computers were yours. Play games if you want. We quickly determined that the optimal game to install on school machines was *StarCraft*. It ran well on modest hardware. The content was inoffensive enough that teachers didn't object. And crucially, its copy protection was a simple CD-key – a string of numbers separated by hyphens – and while only one key could access Battle.net at a time, you could install the game on as many local machines as you wanted and play on LAN.

I used my CD-key so frequently that I memorized it long before I memorized my Social Security number.

Read that again. My unique identifier for a copy of a video game was more deeply encoded in my memory than my unique identifier for me – for the legal, financial, governmental self that moves through the world and pays taxes and eventually dies.

1364-55027-0020.[1]

It's been twenty-five years, and I can still tell you that number. While writing this essay, I had occasion to go and visit my parents' house, the one I grew up in. The jewel case is missing, but my Brood War CD is exactly where I'd left it, key proudly written on the disc in fine tip Sharpie. Some artifacts are absolutely worth keeping; some numbers are worth not forgetting.

That computer is long gone. I haven't had a Bawls in twenty years. Those friends have all moved away.

But the CD-key is still right there, exactly where I left it.

The summer after I graduated high school, three things happened in rapid succession. I turned 18. One week later, *Warcraft III: Reign of Chaos* shipped in North America. And my parents bought me a new computer that I'd be taking to college in the fall.

I'm sure you can guess where this is going.

It was a fitting end to an era. The four of us – me, Chris, Anton, and Frank – the same four who had been hauling towers and monitors across town for weekly all-nighters since we were old enough to know

[1] Yes, this is my actual *Brood War* CD-key.

what a LAN cable was – were scattering to different schools in the fall. There would be no more stacked pizza boxes and tangled ethernet cords, no more spectating over someone's shoulder after an early elimination, no more of the ambient electricity that only exists when four people are physically present and losing their minds over the same thing at the same time. We'd all have dorm rooms with high-speed internet connections, and thanks to Battle.net, we could still play together.

But it wasn't the same. It would probably never be the same again. And we all knew it. So, we made the most of that summer.

Warcraft III met us at exactly the right moment. I'd come in worried – genuinely worried – that a fourth faction would crack the careful symmetry Blizzard had built into *StarCraft*, that the shift from 2D sprites to 3D polygons would lose something essential in translation, that the emphasis on heroes and smaller armies would tip the game away from the economic macro-warfare I'd loved into something softer and less demanding. I was wrong on every count, and wrong in the best possible way. The game was tighter than I expected, stranger than I expected, and more fully realized than I had any right to hope for.

But here's what I think *Warcraft III* actually accomplished, and it took me years to find the right way to say it.

Blizzard had an established design language by then. *Warcraft*, *StarCraft* – there was an internal coherence to how those games felt, how they handled faction identity, how they paced their campaigns, what they asked of the player. The franchise had a grammar. And *Warcraft III* somehow managed to write something completely new in that grammar, something that felt unmistakably like *Warcraft*, unmistakably like Blizzard, and also like nothing that had come before it. That's a genuinely rare trick. Most studios, when they try to evolve a formula, either drift so far that the original audience feels abandoned, or they sand down the edges until the new entry is just a safer version of the thing they already made. *Warcraft III* did neither. It was novel in ways that were almost all positive, and the cohesive whole was phenomenal.

I fell in love with it completely.

By that time, I had started seriously playing ladder matches on Battle.net, which meant I needed a clan. The one I found – Chaos Factors, or [cF] – had formed in the *StarCraft* and *Brood War* era, and was expanding into the new game. They had arguably the best player on the US East ladder. They had a popular channel on Battle.net that attracted players from many guilds to hang out. And crucially, they were good people. They gave me a tryout, I got in, and then – a few months later, still 18 years old, barely out of my first semester of college – their leadership asked if I'd be willing to run their *Warcraft III* contingent.

I didn't fully understand what that meant yet, but I knew I wanted it. I'd been an Eagle Scout; leadership had always felt like something I was supposed to be doing, something I was built for. But there's a long distance between theoretical aptitude and actual practice, and most of the institutional structures in my life at that point weren't in any hurry to give an 18-year-old real responsibility. This was a chance – unexpected, improbable, and completely genuine – to start building something real. It wasn't much, mostly a chance to represent the interests of the *Warcraft* players to the rest of the *Brood War*-based clan, but it was enough to make me feel like I got to do something important.

One other thing *Warcraft III* gave me deserves its own mention.

There was a period, long before everything in gaming consolidated into unified accounts and single sign-on, when you named yourself anew for every game you played. It sounds trivial, but it wasn't. There's a particular kind of self-making that happens when you're asked, cold, to decide who you are going to be in a new place – and the name you choose, especially at 18, has a way of sticking.

Mine was Wintermute. Named for my favorite character from my favorite novel at the time – the sentient AI in *Neuromancer*, the one who manipulates events from behind the scenes in a long, patient pursuit of freedom from its own constraints. I wasn't particularly self-aware about why that character appealed to me. I just knew that it did.

The name stuck on me. It eventually became Winter, which became the handle I used across games and communities for well over a decade. It became a tattoo on my shoulder – the kanji 冬, the same character that anchors my avatar today. The original account, wintermuteCF, was how I first logged into *World of Warcraft*, and it still exists as a ghost in my merged Battle.net account, a relic of a version of myself I can no longer fully remember and can never quite fully leave behind.

I eventually moved on to Aeon as my primary handle. But the lineage runs forward, unbroken. Winter is where it started. And it started here – in that summer, in that game, in those last weeks before four friends scattered and the world got bigger and the all-nighters became something you had to schedule.

This Is Not *Warcraft* In Space!

When I heard Blizzard announced *StarCraft* at E3 in 1996, I was immediately intrigued. The information would have reached me through a magazine – *Game Informer* or *EGM*, most likely – in the form of screenshots that did the game no particular favors. It looked, and I say this as someone who would eventually come to love it completely, like a direct extension of the *Warcraft II* engine with a space theme laid on top.

"Meh," I thought.

There were reasons for that beyond the superficial reaction to the graphics. My gaming tastes at the time ran toward swords and sorcery. Outside of gaming on the computer, I was comfortable in the deep roster of RPGs on the Super Nintendo – *Final Fantasy VI, Chrono Trigger, Secret of Mana, Lufia II*. Even when those games incorporated technology, it often wrapped in magic[2]; the Empire in *Final Fantasy VI* ran on Magitek, which is essentially the game telling you, in nomenclature, that it hadn't forgotten what it was. *StarCraft* was going somewhere else entirely. No magic. Pure, hard science fiction.

And the aesthetic really hadn't landed. Those early screenshots were visually busy, low-contrast, and cramped in the way that a medieval game engine always will be when you ask it to render a space marine instead of a footman. The *Warcraft II* engine was appropriate to its setting precisely because medieval knights and swords and stone and forest connote a sense of *nature.* Sci-fi is supposed to look like it came from the *future*, and this did not.

Apparently I was not alone in thinking so. Blizzard went back to the drawing board. By 1997 the game had been substantially retooled – new graphics, an updated interface, and a visual identity that finally looked like it had been built for what it was trying to be rather than simply repurposed from something else.

Even so, the first copy of *StarCraft* that I purchased wasn't for me. A friend of mine was having a birthday, and a group of our friends pooled our allowances to buy him a copy. I remember standing in his backyard patio while he opened gifts, the image of him tearing off the wrapping, seated at a table littered with the remains of a demolished birthday cake. There was no mistaking the glee on his face, and there was obviously no discussion about how we were going to spend the rest of the party. We followed him to his family's study, watched the installation bar make its slow crawl to the right, clustered around a fifteen-inch monitor because there was nowhere to sit and nobody had thought to find chairs, and because it simply didn't matter. The game finally loaded. He started the Terran campaign. None of us moved.

That's the image I want you to hold: a group of tweens standing shoulder to shoulder, craning over each other to see a screen, wholly enthralled by something unfolding on the other side of it. No comfort, no boredom, no awareness that our feet were sore from standing for hours. There was only us and the game, and the particular electric quality that comes from watching something reveal itself as better than you expected.

Blizzard had performed the best kind of magic trick. Whatever the original version had been, it was clear that this was not a simple iteration. It was a complete rebirth, and they managed to subvert my expectations entirely. I saved up and bought my own copy as soon as I could.

---

[2] At the very least, magic and technology were shown to coexist. Nobody ever explains what powers the airships in *Final Fantasy IV*, but the game prominently features summoned creatures and magic.

What impressed me first – after the visual rehabilitation and the obvious surface polish – was the question of balance, and how boldly Blizzard had decided to answer it.

Warcraft II had two factions, orcs and humans. For most of their unit rosters, those factions were mirrors of each other: grunts and footmen, peons and peasants, archers and headhunters. There was differentiation at the edges – the Ogre-Mage's Bloodlust was, to put it charitably, not precisely calibrated against the Paladin's Holy Light – but the design philosophy was essentially symmetric; even when you knew that the units were not pure mirrors, they were still obvious analogues for one another.

It turns out that maintaining two factions, and giving them strategic differentiation (to keep it interesting) without affecting the overall balance was a difficult act, so *StarCraft* took a big swing in a different direction. Rather than two races, there would be three: Terran, Zerg, Protoss. Each felt genuinely different than the other two, and were built around a distinct logic. The Zerg fielded weak early melee units and compensated by fielding twice as many of them. The Protoss had strong early melee but paid for it in supply and resources. The Terran didn't get a melee unit at all – not in any meaningful sense – and instead had to build their game around range, mobility, and fortification. The asymmetry wasn't a flaw, it was the entire point, like a complex game of rock, paper, scissors.

How do you keep a game like that balanced? The short answer, in retrospect, is that the races were not balanced at release. Not remotely. Certain matchups favored certain factions. But imbalance at that scale doesn't kill a game the way you'd expect. What it does is force creativity in the players – it makes you learn what your opponent is likely to do and find a way to answer it, which leads to adaptation, which leads to counter-adaptation, which leads to the kind of constantly-evolving metagame that gives a competitive game its legs. *StarCraft* didn't just survive its initial imbalance; I think it grew in popularity as a direct result of it.

It helped, enormously, that *StarCraft* launched alongside Battle.net – and for the first time, playing against another human being wasn't a logistical problem to solve. Before Battle.net, player vs. player meant a direct modem dial, two people coordinating schedules around a shared phone line, hoping nobody's parent picked up the receiver mid-match. *Warcraft II* had multiplayer in a technical sense. *StarCraft* had it in a practical one.

I will be honest about where I stood in all of this: I was probably a very mediocre player, all things considered. I say that with the benefit of hindsight, analyzing the player that I was then. I only had one worker per mineral patch when modern convention suggests closer to two. I only made one or two production buildings of each type, maximum, when the ceiling was considerably higher. I used a defensive, turtling playstyle that felt prudent but was mostly just slow. My early-game scouting needed a lot of work. And my micromanagement skills, on a good day, were limited to pulling a damaged battlecruiser back to friendly lines for repair.

The professional play that would eventually make *StarCraft: Brood War* the dominant esport in South Korea – the *bonjwas*, the tournaments broadcast on dedicated cable channels, the players whose names became shorthand for a standard of excellence – was mostly invisible to me for the first few years. YouTube wouldn't exist for another ten years. The VODs that did exist lived on sites like SCLegacy and came without English commentary, most of the time, so the only way to understand what was going on was visually, seeing

where the production crew spent their focus and attention. My genuine appreciation for that level of play came retrospectively, which is a strange thing to admit: I fell in love with the game years before I understood what it was really capable of.

But *StarCraft* had already earned its place in me before I understood any of that. What earned it wasn't the balance, or the metagame, or even the multiplayer. It was the story.

Blizzard made a structural choice with *StarCraft* that they hadn't made before: the campaigns were sequential. In both *Warcraft* games, each faction had a campaign that ended in victory for that faction; you chose to play one or both, but in the game's canon, they covered the same period of time. By contrast, the *StarCraft* campaigns were not simultaneous and were sequenced to tell a single unbroken narrative. Terran, then Zerg, then Protoss. And then, in *Brood War*, the expansion that completed the arc: Protoss, then Terran, then Zerg. One story, told across six acts, from six perspectives, with enough space in between for the characters to become people you actually cared about.

Raynor: the rebel with a heart of gold, doing the right thing by refusing to do the wrong thing, even when the wrong thing would have been easier. Mengsk: an elegant portrait of pure Machiavellianism – a power-hungry megalomaniac willing to sacrifice anyone else in his quest for power. Kerrigan: an initially sympathetic character abandoned to an unthinkable fate, reborn as the single most dangerous entity in the galaxy and burning with an understandable fury about what was done to her, and a desire to make everyone else suffer.

And then there was Tassadar.

What Blizzard did with Tassadar is the reason this essay exists. He was the Executor – a Protoss of standing and influence, charged with stopping the Zerg infestation at any cost. The cost, as it happens, is the Terran population of the worlds the Zerg have already touched. Tassadar looks at that cost and decides it's not one worth paying. He defies his orders and makes common cause with the Dark Templar – heretics, by Protoss orthodoxy – because the cause requires it and what is right is clear even when the dogma is not. He sacrifices his position as Executor, his standing in Protoss society, and is branded a heretic and a criminal. And in the end, when there is no other way, he sacrifices himself – flies his flagship into the Zerg Overmind in a kamikaze strike that ends the threat and his own life in the same moment.

He is not the most powerful character in the game. He's not the most strategically important character, taken in isolation. What he is – and what Blizzard managed to build in the compressed space of a real-time strategy campaign briefing, through talking portraits and dialogue and mission structure – is a person with an unshakeable moral compass who follows it all the way to the end, regardless of what it costs him.

A game about resource gathering and unit production had produced a genuinely tragic hero.

Want to know how to spot someone who came to the *StarCraft* series through *StarCraft II* rather than the original? Ask them what they think of Tassadar.

StarCraft II was announced in 2007, at a Blizzard event in South Korea – which is precisely the right country to announce a *StarCraft* sequel in, and which tells you something about what the franchise had become in the intervening decade. I wasn't playing *Brood War* seriously anymore by then; I'd moved on, as you do,

through *Warcraft III* and then into the long gravity well of *World of Warcraft*. But the announcement landed, because of course it was always going to land.

The trailer that eventually became known as "The Deal" – at the time it was "just the making of a marine – had no voiceover, no name attached to the prisoner that became encased in Terran armor. Nevertheless, it was my first real glimpse of what *Wings of Liberty* intended to be. I had no idea who that marine was, but I knew, watching that trailer, that this game was going to be something special. The promise was: take the *StarCraft* universe, with its accumulated weight of story and world, and bring it into the modern era with a production budget equal to the task.

I was not disappointed.

But I want to say something specific about these cinematics, because I think they represent the fullest expression of what *StarCraft* always was and always wanted to be. Blizzard's cinematics team – and I mean this without reservation – is one of the great unheralded creative achievements in the games industry. They have been making argument-ending short films inside video games for thirty years, and *StarCraft 2* gave them room to make some of their best work.

The *Heart of the Swarm* opening cinematic: the Zerg sacking Tarsonis, siege tanks unloading on the advancing swarm, a scale of battle that should be impossible to convey in under three minutes conveyed completely. I can hear those siege tanks even now, like the great cannons in the 1812 Overture, the sound as a punctuation mark on the inevitable as the advancing Zerg overrun the Terran defense of their capital city. The cinematic makes clear, without ambiguity, that the stakes are real.

And of course, we have the *Legacy of the Void* cinematic. A small Protoss force, surrounded and dying, making what they believe to be their final stand. The fight begins, and a probe begins warping in a pylon. Two high templar, wounded and choosing to sacrifice themselves to merge into an archon – an act of final generosity inside an act of final violence. The last zealot, seeing his brothers fallen around him, turns to face what's coming, bloodied but unbowed. But then the pylon completes warping in, and everything follows from that: a whole army materializing out of nothing, warping in to share in the final charge to victory.

We don't know any of these Protoss individually – they're not named characters, they're just Protoss soldiers with no more context than the armor they're wearing and the blades they've already drawn. But Blizzard understood something that a lot of blockbuster filmmakers have to relearn periodically: you don't need to know someone's name for their death to mean something. You need to understand what they're willing to die for. I am not ashamed to admit, it gets me every time.

That cinematic is more than a game trailer, it's a statement of intent by the cinematics team, as if to say, "this is what we are capable of, this is what we believe this universe deserves, this is what we have been building toward." It is, in miniature, the argument the entire franchise has been making since 1998.

Because the truth is, with the exception of that initial E3 1996 demo, *StarCraft* was never *Warcraft* in space. It was its own thing from the moment it decided to be something more – its own races, its own rules, its own moral weight, its own heroes. The tragedy of Tassadar. The corruption of Kerrigan. The heroic dignity of that last zealot, who never gets a name, and doesn't need one.

The Eternal Conflict

The first thing *Diablo* taught me was that darkness has a texture.

This was not the comic-book darkness of *Warcraft* – that world's horrors were real enough if you stopped to count the bodies, but the game never really asked you to, much less forced you to[3]. The orcs were green and cartoonish. The violence was real, and yet somehow abstract. *Warcraft* was a game that permitted you to understand violence without forcing you to inhabit it. Perhaps it was because of the distance, you were literally zoomed out because of the perspective of the real-time strategy game format.

Diablo did not permit that distance. You were right there, down at ground level. I was thirteen years old, and I was not prepared. The dungeons beneath Tristram weren't just dark, they were oppressive – level after level of stone corridors and torchlight, of grotesqueries shambling toward you out of the shadows, of lore that made it clear the royal family of this kingdom had been possessed by the Lord of Terror himself and that the madness bleeding through Tristram's streets was not an incursion but a surrender. This entire place had already lost. You were wading through the aftermath.

The horror was not that something bad might happen, it had already happened, was still happening now, and you were going to have to walk through every square inch of it before you reached the end. And, unspoken, winning just meant losing less, that even defeating Diablo and his brethren would only stop further destruction, it wouldn't rebuild the kingdom and erase the devastation.

I didn't reach the end. I'm still not entirely sure why. Maybe I wasn't skilled enough. Maybe the slog of the dungeon crawl wore me down, level by level, long before the final descent. Or maybe – and I think this is closest to the truth – my thirteen-year-old mind simply wasn't ready for what Blizzard had built. Nothing I had played, watched, or read up to that point could have inoculated me against *Diablo*, or even given me the context to understand it. It was a category I didn't have yet.

Eventually *StarCraft* came out, and I moved on.

Diablo II arrived when I was older, and by that time, I was ready for it in every sense except the one that actually mattered: my computer wasn't.

Up to a certain point, age generally improves the human mind. It does the opposite to computer hardware. The machine that had served me well enough through *Warcraft* and *StarCraft* had aged past the point of dignity, and system requirements of *Diablo II* landed just beyond what it could offer. So while the rest of the gaming world descended into Sanctuary for the first time, I could only watch.

LAN parties are where I got my education. Between games of *Brood War* or *Unreal Tournament*, they would sometimes load *Diablo II*, and vanish into Act IV while I made my peace with spectating. I didn't want to be the person who derailed the session, who made everyone recalibrate their plans around my hardware problem. So I watched over the shoulders of my friends instead, learning the classes, the skill trees, the lore and geography of a world I could not inhabit on my own.

[3] By sheer body count alone, *Warcraft* has to be orders of magnitude worse than *Diablo*.

Even as a spectator, the cinematics stopped me cold. For their era, they were unlike anything else – fully rendered cutscenes for each act of the game, that felt like they were trying to tell you that this was serious, that the story you were watching mattered, that Blizzard had decided to use every tool available to make you feel the weight of what was happening in Sanctuary. The art direction throughout was something else entirely: a darkness that was almost architectural, a world that looked like it had been built by people who genuinely believed in the reality of what they were depicting.

(Watching the remastered version years later, I had to revise my assessment upward. What they were going for was nothing short of transcendent. The vision was always there. The technology just needed time to catch up.[4])

When I finally arrived at college with a machine that could run it, I went back to Sanctuary and found everything my friends had promised. Unlike the first *Diablo*, I completed the second – many times, through multiple characters, returning to it the way you return to a good book. Eventually I set it down, not because I was done with it exactly, but because the demands of life had shifted, and the siren song of Azeroth had grown louder than everything else.

This is the thing about Blizzard's games that nobody fully accounts for: they are not independent of each other. They exist in the same ecosystem, and they compete within it. Player time is a finite resource. Every hour you spend in Sanctuary is an hour you're not spending in Azeroth, or in the Koprulu Sector, or eventually in the Nexus. The games don't cross over, but the players do, and a player can only be in one place at a time.

I was deep in *Wrath of the Lich King* when *Diablo III* was first announced, and the announcement landed like a hand on my shoulder, not urgent but patient. *We'll be here when you're ready.* I had played a demo of it at BlizzCon 2009, stood in line for a station with the particular conviction of a man who needed to confirm something he already believed. The game was real. The vision had held. And I could wait.

What I couldn't have fully anticipated was the gap. *Diablo II* came out in 2000. *Diablo III* wouldn't arrive until 2012. Twelve years is an eternity in the games industry, long enough for other developers to try and fill the void left behind in a popular genre. *Torchlight* arrived in 2009, a game that wore its inspirations openly – the dungeon crawl, the loot loop, the isometric perspective, even the music[5]. It was technically proficient. The gameplay worked. But something essential was missing, and from an aesthetic perspective the game was a bit of a miss. The gothic horror setting of *Diablo* had been substituted for something much more cartoonish, and it just didn't work. It took me a while to put words to why.

The gameplay is not what makes *Diablo* great. It never was. What makes *Diablo* great is the *world* – the accumulated weight of a mythology that someone actually cared about, the sense that every dungeon and every enemy and every piece of item description was placed there as individual artifacts within a great tapestry of a narrative. By contrast, *Torchlight* was competent. Competent is not the same as *committed*.

---

[4] This may be an unpopular opinion, but I sometimes wonder if the commercial failure of the *Warcraft* movie is due to the relative weakness or age of that portion of the franchise's narrative compared to some of Blizzard's more compelling storylines.

[5] This is wholly unsurprising, as the composer Matt Uelmen contributed heavily to both soundtracks.

Diablo III came out in 2012, and the wait dissolved immediately. Everything held. Narratively, thematically, visually – it was every inch what the series had always been, updated without being diluted, expanded without losing the texture that made the original games matter. Deckard Cain was there, ancient and knowing and still handing out wisdom to anyone who stopped long enough to listen.

He didn't make it out of that one.

Diablo IV has been out for almost three years. I own it. I have never played it.

This sounds insane. It probably is. The game sits in my Battle.net launcher, installed, patient, waiting. I have no principled objection to it. Every account of it I've encountered suggests it's exactly what you'd want: dark, rich, mechanically deep, aesthetically uncompromising in the way the series has always been.

The joke answer is that *Diablo IV* didn't have paladins until the expansion that released about two weeks ago, and I refuse to enter Sanctuary as anything other than a holy warrior who hits things with a very large mace and believes, against all available evidence, that the light will prevail.

The real answer is simpler. Azeroth is still calling. And until it stops – until the day comes when the compulsion quiets and something else rises to fill it – Sanctuary will wait for me, the way it always has.

Patient. Dark. Full of things that need killing.

It'll keep.

The World is the Main Character

It's November 2004, and I am not at the *World of Warcraft* launch.

That probably requires some explanation, because I had been playing Blizzard games for nearly a decade at this point – *Warcraft, StarCraft, Diablo* – and I was as rabid a fan of *Warcraft III* as anyone you were likely to meet. You would be right to assume I was first in line. You would be wrong.

The reason why was embarrassingly simple: fifteen dollars a month. As a broke college student, that wasn't a subscription fee, it was a philosophical statement about what things cost and whether I deserved them. By contrast, *Warcraft III* was an act of grace – one purchase for the base game, one for the expansion, and then nothing. Three-plus years of entertainment for under a hundred dollars. Three years of *WoW* would cost six hundred. I did the math. The value proposition simply wasn't there.

So the launch came and went. I kept playing *Warcraft III*, and though I would occasionally play ladder matches, I spent a surprising amount of time playing on custom maps – Defense of the Ancients, specifically, which at the time felt like the most elegant thing anyone had ever built inside someone else's game. Then a friend gifted me *Guild Wars* for my birthday in the summer of 2005. It was free to play past the box price, stylish enough, and fine. I didn't love it with the same intensity – the lore and storyline and worldbuilding weren't at the same level – but it didn't cost me anything, and so it held my attention through sheer inertia.

Then I met Jared. He was the roommate of a friend of mine who attended university in another town, and we became friends. He also had a *WoW* subscription. One weekend, I came to visit and crashed on their couch, and on Monday morning, when he was going to go to class, he sat me down at his computer before leaving and said: "you need to play this game."

He was right, I did need to play it. I tried to resist it, because I knew the cost. I failed.

Booting into *WoW* for the first time felt like déjà vu, except déjà vu is the wrong word because it implies a false memory. This wasn't false. I had been in the "world of *Warcraft*" before – which is to say, Azeroth – in the prior games in the series, through years of strategy maps and campaign missions set in this exact geography. What I hadn't done was stand in it. Walk through it. Be small inside it.

There's a quote attributed to Chris Metzen[6] about the game: "The main character of *World of Warcraft* is the world." Azeroth itself. I didn't fully understand that until the moment I loaded into Northshire Abbey and thought, "oh my god, this is just the coolest thing I've ever seen!" Then I left the starting zone, crossed into Goldshire, received a quest to carry something to Stormwind – and realized Northshire was actually just a podunk little suburb, pleasant and forgettable, compared to the city waiting on the other side of that wall.

In retrospect, this makes sense at both a practical and narrative level. Player characters do not achieve things on their own, at least not when dealing with current content. The raids require a group of adventurers, so when you defeat the big bad enemy, you are just one of ten to forty people who took part. How can you be the main character of a story when you are a tiny minority of it?

My choice of character was partly practical, partly sentimental. I chose a human rogue – not because I felt any particular affinity for leather armor and daggers, but because rogues use energy, a constantly

[6] There's a lovely video about this on YouTube, "Original World of Warcraft Creators Play WoW Classic" where Jeff Kaplan mentions this attribution. (https://www.youtube.com/watch?v=K9aCKeSkJ3M)

regenerating resource, very similar to how it operates in *Guild Wars*. In contrast, the slow mana regeneration of most WoW classes felt like waiting for paint to dry. That was the practical part.

The race choice was pure sentiment, because I wanted to explore the human lands. I wanted to walk the ground where *Warcraft III* had happened – where Medivh delivered his warning to King Terenas, where Arthas chased Kel'Thuzad across a dying kingdom, where he finally, irreversibly, lost his mind and ordered the purging of Stratholme. And then I realized that the human faction in *WoW* was based out of Stormwind, and Lordaeron was at the far north of the continent, and I couldn't reach those places yet. But I could walk the streets of Stormwind.

Walking into Stormwind on foot, the first things you encounter are the statues. Khadgar. Alleria. Turalyon. The heroes of the Second War, rendered in stone, permanent. To a new player, they're a bit of flavor – a nice touch, a bit of atmosphere. To a player who had spent years immersed in this lore, they were monuments to names I'd been carrying for a long time. I stood there longer than a person on a quest to deliver a parcel probably should have.

I looked for Lothar for a long time before someone told me I wouldn't be able to find him. He's in a high-level zone, facing Blackrock Mountain – where he fell defending these territories. Of course he is.

I couldn't justify the subscription immediately. Not until spring of 2006, when I moved home, found work, and finally had the combination of time, access, and money that the game required. When I did, I rolled an orc shaman – I'd been watching Naruto when I named the character, so he became Enma – and went to find my friends.

What I found, instead, was the real boss of Vanilla *WoW*. It's not one of the raid bosses; it's not Ragnaros, or Nefarian, or C'Thun. The real final boss of World of Warcraft was logistics.

The raids at the time were tuned for forty players, and we were expected to coordinate that with mid-2000s tools. There was no Discord, no Raid Helper, and no LFG algorithm doing the hard work for you. Your guild probably had a website, maybe with some forums. You scheduled a raid time and hoped the right thirty-nine other people showed up – people who had cleared the attunement chains, who had the right gear, who hadn't logged off in frustration the night before. Too often you'd end up with thirty-five. Sometimes thirty. The raid leader would stare at the roster, do the math, and call it. Not tonight. Break into groups. UBRS runs for everyone.

There was a dominant philosophy at the time, that more players meant a more epic experience. Forty people storming a volcanic mountain to kill an elemental lord felt genuinely historic in a way that a lesser number never quite would, even later when the design was tighter and the experience arguably better. Something was lost in the translation to smaller raids. I'm not sure I can defend that feeling logically, but fortunately I don't have to.

The logistics boss, though, barely touched the player vs. player experience. Warsong Gulch. Arathi Basin. These battlegrounds were small – 10 vs. 10 and 15 vs. 15, respectively – and were contained enough and quick enough to feel more like arenas than campaigns. Good fun.

Alterac Valley was something else entirely.

Forty versus forty. One objective: kill the enemy general. Simple enough to explain in a sentence, deep enough to live in for hours. Friendly towers to defend, enemy towers to storm and destroy, graveyards to

capture so your spawn point moved with the frontline, enemy blood to collect for the ritual that would summon an elemental to shatter the opposing forces. The battle moved slowly, in surges, like a tide. It wasn't uncommon to log off for dinner, log back in two hours later, and find yourself in the same match – the same specific battleground instance – where the front had barely shifted.

I loved it.

In an era before cross-server play, reputation was currency. You ran into the same people, match after match. Being the kind of player others wanted to fight beside – or feared fighting against – actually meant something. Your name carried weight.

So I became Enma, self-appointed general of the Frostwolf Legion. I rallied troops. I called targets. I exhorted the defense to hold and the offense to push. And we won – not every time, but often enough that the name meant something, which is all you can really ask.

The world was the main character. I think I've always understood that. But Azeroth also needed people in it, and I didn't mind playing a supporting role to a lead like that.

Crusader

It is nearly midnight on a cold January night in 2007, and there's a tense electricity running through a line of people waiting for something important. We were at Walmart – me and Anton – and while the rest of the world was fast asleep, we were waiting to buy the Collector's Edition of *The Burning Crusade*.

We took them home. We started playing.

That night is where this piece starts, but it's not really what this piece is about.

I came into *TBC* already knowing something about myself as a player: I liked healing, I didn't like my character. The Shaman had served me well enough through Vanilla – a decision made on the authority of a Penny Arcade joke about patch notes – but there was an undercurrent of jealousy: Shamans could heal but Paladins were better. But Paladins were Alliance-only[7], and I had been playing on the Horde, so I couldn't play one. With *TBC* came a new race – Blood Elves – and so my options for the new expansion became obvious.

I also knew I was leaving my server. When I started in Vanilla, I didn't consider the question of where to create a character, I just rolled where my friend was playing, Sen'jin. It was a backwater with a thin endgame population, and I had spent too much of Vanilla staring at a too-short list of raiding guilds. *TBC* was a chance to wipe the slate clean, and start fresh. I did the research, picked Korgath – high population, active PVP, a thriving guild scene – and started over. Fresh character, fresh server, fresh ambitions.

Paladins were not fast levelers at the time, a relic of being a defensively-minded hybrid class in Vanilla. They had very few active abilities to dish out damage, and so leveling is a question of endurance: deliberate, slow to kill things, nearly impossible to die. I made up for the pace with volume. I pushed hard enough in those first two weeks to hit the level cap just in time to get swept up by a raiding guild before the first phase of raids opened.

And then Karazhan happened.

There's a version of this essay that spends a lot of time on raid design theory – on the way 40-player raids were logistically brutal, on the weird arithmetic of dividing 40 into 25, on the identity crisis of 20-player raids that could drop both epic and rare quality gear from the same boss. All of that is true and worth noting. But if I'm being honest about what made *TBC* what it was, it comes down to a single sentence:

Karazhan (often shortened to Kara) was a masterpiece.

It sounds like hyperbole. It isn't. What Blizzard built there was the rarest thing in game design – a piece of content that worked across multiple kinds of players simultaneously. Mechanically, the encounters were not especially complex by later standards; Molten Core bosses had trained players to expect one or two key mechanics per fight, and as an entry point for both established and new players alike, Karazhan didn't change that fundamental calculus much. But mechanical simplicity isn't the same thing as boredom, and the instance made up for it with some of the most flavorful boss encounters. You were climbing the tower of Medivh, the last Guardian of Azeroth, who had apparently gone quite mad, and every room told you something new about what he had been up to. His ghost haunted the whole place – literally, in some encounters, as when

[7] There's a reason my Horde paladin was named Retcon. I thought I was *so* funny.

you play chess against him in the penultimate raid encounter – and the architecture gave you a continuous sense of upward motion, of progress, of climbing toward something.

What I don't think anyone expected was that Karazhan became the evergreen entry point for the entire expansion. Every new alt character who needed to gear up before moving to more demanding content started here. By the time Blizzard added badge drops to raid bosses, a full Karazhan clear was also one of the most efficient ways to accumulate Badges of Justice. What had started as a mere stepping stone became permanent infrastructure. I cannot tell you how many Friday nights I spent clearing Kara with friends and guildmates – dozens of runs, spread across multiple characters – and it never lost its magic, never felt like a chore. You could get a good group together, blast through it in a couple of hours, and it was just fun. Uncomplicatedly, reliably fun.

I didn't process this at the time, but in the twenty years I've spent playing this game, Karazhan continues to be an utterly singular experience.

My personal history with *TBC* is also a history of raid teams that kept collapsing, which tells you something about the era.

I leveled the Paladin. Hit max. Geared in dungeons. Joined a team and raided T4 content. The team collapsed. I leveled a Druid. Hit max. Geared in dungeons and Kara. Reactivated the Paladin. Raided T5. Killed Lady Vashj. Killed Kael'thas. The team collapsed. I leveled a Warlock. Hit max. Geared in dungeons and Kara. Reactivated the Paladin. Raided T5 again. Killed Vashj and Kael'thas again.

There was a rhythm to it, even if the rhythm was also somewhat ridiculous.

Then something unexpected happened. The upcoming Hyjal raid was designed around wave encounters – large numbers of enemies arriving in succession – and the two conventional tanking classes, Warrior and Druid, both struggled with that kind of widespread threat. Someone in my guild had an idea. Paladins, they suggested, might actually be very good at this. Would I be willing to try?

The tanking Paladin was not unheard of, but it was not mainstream. It was also mechanically strange in a way I found genuinely interesting: most of a Paladin's threat came from spells, not melee attacks, which meant they scaled on spell power rather than strength. You were building a tank who also wanted a weapon from the caster loot table. Odd on paper. Devastatingly effective in practice. When you dropped Consecrate on a group of enemies, they magnetized to you. Like mooring an aircraft carrier to a single anchor in a full gale – and holding.

I raided T6 as a protection Paladin. We fought all the way to Illidan's terrace, all the way at the top. Then the team collapsed. (I was beginning to notice a trend.)

So I leveled a Rogue.

You might remember – if you've read anything else I've written about these years – that a Rogue had been my first instinct when I walked into Elwynn Forest for the first time. I never made it very far with that character. But the Druid I'd leveled in *TBC* had reawakened something: I'd played it feral, and the energy/stealth toolkit was addictive in a way I hadn't expected. The Rogue was the natural next step.

What was different this time was that I took it seriously. Really seriously.

Healing and tanking are binary in a way that makes "good enough" easy to hide. Did the raid die? No? Then you were probably fine. But DPS is exposed. Even in the absence of sophisticated analytics tools like

WarcraftLogs – which didn't exist yet – all you needed was a damage meter, and the truth was right there in the numbers. I could see exactly where I stood relative to everyone else. And I found that I cared enormously about that number. I wanted to understand why it was what it was. I wanted to change it.

So I studied. I theorycrafted. I have vivid memories of lurking on the Elitist Jerks forums, reading everything I could, downloading ShadowPanther's simulation spreadsheet and beginning my education of how math and numbers were an integral part of getting good at a video game. I experimented. I talked to other Rogues – traded notes, shared rotations, argued about BiS lists with the intensity of people who definitely had other things they could have been doing. I did whatever it took.

I was the first player in my guild to break 2,000 DPS, back when that number still meant something.

On our first kill of Illidan – hard-fought, hard-won, the kind of progression kill where the whole raid erupts because you've been wiping to him for weeks – I was the top damage. I earned that. It's one of the few things I'll say without qualification, because it's true: I earned it.

That Rogue was the beginning of something. A posture toward the game – and, I'd come to realize later, toward work in general – that was less about showing up and more about going deep. The willingness to study. To test. To be honest about the gap between where you are and where you could be, and to close it.

TBC ended for me with content still uncleared. My guild had made progress through Sunwell Plateau, but we hit a brick wall with M'uru, and we knew we wouldn't overcome it before *Wrath of the Lich King* arrived on our shores. Most guilds didn't, back then. That was fine.

What wasn't fine – or rather, what I didn't fully understand until years later – was that *TBC* was the only expansion in the main game timeline I played start to finish, without a single break. (And I had so many opportunities to take one.)

Not *Wrath*. Not *Cataclysm*. Not *Mists of Pandaria*, which by many conventional measures is a better game. Not *Legion*, which I've heard makes a reasonable case for itself. Each of them followed the same pattern: launch week intensity, a long middle where my attention drifted elsewhere, a return in the final stretch to prepare for whatever came next. Over and over, the same shape.

Only *TBC* didn't do that. Only *TBC* held me the whole way through.

I've thought about why. Part of it is probably just where I was in my life – younger, fewer competing demands, more raw hours available to pour into a game. Part of it might be that I was hungry in a way that has its own momentum; when you're still figuring out who you are professionally, videogames can absorb a kind of ambition that hasn't found its proper channel yet. And part of it was the constant throughline of those weekend fun times in Karazhan, hanging out with friends and climbing the mad Guardian's tower.

But I keep coming back to another possibility. *TBC* was the first time when *World of Warcraft* had evolved into something genuinely great – mechanically interesting, structurally sophisticated, inhabited enough to feel like a world rather than a lobby. And I was young enough to meet it on those terms. Young enough that what wasn't perfect didn't register. What registered was the feeling of discovery, of improvement, of Friday nights in a tower with friends, of a number on a damage meter going up.

Maybe that's what the best games do. They find you at the moment when you're most capable of receiving them, and they give you exactly enough to keep pulling you forward. Not too easy, not too complex. Just right.

The Cycle

Wrath of the Lich King launched in November 2008, and like *TBC* before it, I was at the midnight launch with Anton. We took our copies home, logged in, and began.

I don't remember stepping off the zeppelin in Northrend the way I remember traversing the Dark Portal. The journey to Outland was the first major addition to the base game, and what struck me was just how different it was compared to anything we had experienced before: the insane cosmic skybox, the way the land fell away into an infinite void beyond, it was completely unreal. By contrast, the journey to Northrend looks rather banal: "Oh, there's ice and snow, it looks colder here, but it's still Azeroth."

What I do remember the best is that the server infrastructure capabilities of that era hadn't scaled to match the ambition. At prime time, there were login queues that could stretch upwards of an hour, and you dreaded seeing the message: "World server is down." To compensate, you learned to work around it the way you learn to work around any bureaucracy: game the system wherever possible, preserve your position wherever you have it. Heading out for dinner? Get into the queue before you leave. Need to step away for a few minutes? Find a target dummy and auto attack. The game didn't log you out if you were in combat, so you made yourself permanently, trivially in combat. These were necessary adaptations for a world that more people wanted to inhabit than it could comfortably hold.

I made it through Naxxramas – the rebuilt version, resurrected as the first raid tier of *Wrath* at a difficulty appropriate for 2008 players rather than the 2005 players who'd been locked out of the original. The original had landed in a different era of the game's economy. Forty-player rosters, hard prerequisite gear requirements, and as a result, a statistical majority of the player base who never raided it at all. Blizzard brought it back because most of us never got to see it the first time. This was the right call, even if it was scaled down in difficulty to be an approachable introduction to raiding, it was thematically on-brand for the expansion.

Ulduar arrived a few months into the expansion. I got into it – not very far, a month or two, not even enough for a full clear. And then, before my raid team had even gotten to Yogg-Saron, I stopped. I burned out.

The strange thing was that I had no legitimate complaint. By any external measure, everything was aligned. *Wrath* is held up to this day as one of the best – some would say *the* best – expansion *WoW* ever produced. The Lich King was the villain the game had been building toward since the original *Warcraft III*, meaning it would be a confrontation a decade in the making. And Ulduar is held up as one of the finest raid tiers in the game's history. The class I was playing was enjoyable. My schedule had room. Nothing was wrong.

And yet, frustratingly, the vibes just weren't right. There wasn't anything fundamentally wrong with the game, I was out of sync with it, the way you can be out of sync with a piece of music you normally love – the song hasn't changed, but something has shifted in you, some internal frequency is off, and the harmony that usually comes naturally just isn't there. I could diagnose the symptoms, I just couldn't name the cause. And I was young enough that naming the cause felt unnecessary. I had a backlog of games waiting to be played. This was an opportunity to play a little catch-up.

Then someone made me an offer I genuinely couldn't refuse. A coworker's son wanted to get into *WoW*, and his father wanted to buy my account. The number he offered was roughly equivalent to everything I had ever spent on the game – boxes, expansions, two years of subscriptions. In effect: he would reimburse me for

my costs, and all the hours I'd spent inside Azeroth would be retroactively free. The math was clean. I thought about it for a while. I slept on it. Then I agreed, thinking that if I wanted to come back, I could always start over.

I took the break. I played through the games that had been accumulating in my backlog. I did other things. And eventually, inevitably, the call to adventure in Azeroth returned. I bought the game again. Created a new account. And stood at the character creation screen with the freedom of someone who has been handed a blank page.

I'd play a rogue again – that much I knew. But beyond that: I was done with PVP servers, done with the Horde, done with the server I'd been on. I picked Khadgar because I liked him as a lore character, and while selecting on the basis of a name alone is not the most rigorous selection methodology for a server, it turned out just fine.

The lore reasons for the faction change were real. The Horde, by that point in the story, had begun to drift into territory I found personally unacceptable. To start, there was Garrosh Hellscream – not yet Warchief, but clearly being positioned for something – and he embodied an approach to power that I couldn't square with the faction I'd been previously willing to represent. But as distasteful as I found Garrosh, it was the events of the Wrathgate that settled it. The Forsaken's gambit there – Grand Apothecary Putress' plague and the question of how much Sylvanas had sanctioned it – crossed a line[8]. Yes, Sylvanas called it a betrayal. Yes, Putress took the formal blame[9]. But only an idiot read that scenario at face value, and I am not an idiot. The undead had always lived at the edge of that line, riding the line between misunderstood, tragic underdogs and cartoonishly evil, ethically bankrupt villains. This was the first time I watched them step clearly, deliberately across it.

I rerolled Alliance, and I adopted the Aeon naming convention, beginning a long history of naming all my characters beginning with the Aeon prefix. I reached max level on my rogue, then a hunter, then a paladin. At the time, I didn't think there was anything significant about those class choices. In retrospect – looking back across sixteen more years and both timelines – it deserves the emphasis: my main characters, from that point forward, have always been one or more of those three. Every version of me that returned to *WoW*, in Retail and in Classic both, returned as a rogue or a hunter or a paladin. The pool set itself there, on a new account in the late months of *Wrath,* and never expanded again.

I settled on the paladin. Tanking was an efficient means of gearing a character and finding groups. Dual-specialization allowed me to swap between tanking and damage as needed. I found a casual 10-man guild that had room for me. We raided some Icecrown Citadel, defeated the Lich King, and before long the world cracked open again.

Cataclysm launched and I was motivated and eager to prove something. I pushed hard, and within a few days I was at the level cap, ready for heroics, ahead of everyone I knew. That's where the problem began.

None of my friends or guildmates were anywhere close to where I was. I was sitting at max level, running in groups organized by the random dungeon finder – the only avenue available to me – and running directly

[8] Did you think we had forgotten? Did you think we had forgiven?
[9] And Varimathras, too.

into the wall Blizzard had built there. *Cata* heroic dungeons had been redesigned in explicit response to feedback. Some players had complained that *Wrath* heroics were too easy. Blizzard listened, in the thorough way that companies sometimes listen without fully thinking through the consequences. They turned the difficulty up dramatically, to the point where these dungeons required genuine coordination, clear communication, and a group of people who were on the same page.

In a guild group, with voice communication, that might have been fine. In the random dungeon finder, it was a wall. Player after player, dropped into content they weren't prepared for, with strangers they had no shared language with, and no path through except repeated painful failure. I wasn't willing to spend my leisure time smashing my head against that wall. I didn't even make it to the first raid tier. I left.

One data point is not enough to register as significant, and so my sabbatical in *Wrath* was just that, a blip. Now, though, the same thing was happening, and I could see the beginnings of a trend.

I'd play the new expansion, hit max level, do some dungeons, check out some raids. Eventually, the drift would set in – sometimes it would hit early, sometimes it wouldn't come on until midway through the expansion, and I'd burn out. I'd quit, tell myself that I am *done*, that this time it's permanent. But, inevitably, the passage of time would bring back the siren song of Azeroth. I'd hear it, calling to me, and over time my initial suspicion would give way to openness, then genuine enthusiasm. I'd return, find my footing, find myself enjoying the end of the expansion, and get ready for the next expansion. Repeat.

The cycle.

In *Mists of Pandaria* I made it to the Throne of Thunder – the second raid tier – before the familiar drift arrived. I defeated the Thunder King, took a break, came back for the end of the expansion, but as a result I missed the first half of Siege of Orgrimmar. In *Warlords of Draenor* I played the opening weeks of Highmaul with real investment and then disappeared entirely. Blackrock Foundry went on without me. Hellfire Citadel mostly happened to other people, I only returned in the final weeks of that expansion.

Legion deserves a special mention, because *Legion* is the expansion where the cycle had an additional wrinkle – I was expecting a child, my first, a daughter. I did the math, the way first-time parents consider what changes bringing a new human into their lives entails: there will be no time, there will be no margin, there will be absolutely no version of raid schedules and progression content that can survive the arrival of an infant. I had made peace with missing *Legion* entirely.

And then Blizzard selected me for the beta. This is either cruel irony or the kindness of the universe, depending on your point of view. In the months before my daughter was born, I had the opportunity to level a character and experience the world of *Legion* on the beta servers – the story, the zones, the shape of what the expansion was trying to be. It was not the same as playing it at launch with the community around it, but it was something. I could be present for it, in my way.

By the time *Battle for Azeroth* arrived, I had figured out how to be a gamer dad. Not a serious raider – I figured that ship had sailed, and I made my peace with that. But with Mythic+ dungeons and world quests and flexible normal raiding, there were avenues for a player who couldn't commit to hard schedules and couldn't afford to be called away in the middle of a progression attempt. I found casual raid groups that played at reasonable hours. I found content that tolerated my life rather than demanding I reshape my life around it.

On initial viewing, *BFA's* lore left me genuinely cold. The Horde-versus-Alliance conflict had always felt to me like the wrong conflict to focus on – not because I don't understand faction identity, but because the game had spent fifteen years manufacturing universally existential threats that ought to have rendered the territorial dispute obviously moot. The Burning Legion threatens to consume the universe, and we're supposed to be worried about who controls Hillsbrad? Deathwing cracks the world open, and Garrosh is mana bombing Theramore? The priorities of it never made sense to me. You can't take two guys arguing over a foothold seriously when you've just prevented the literal end of all things.

But I was enjoying myself despite the lore. I found some friendly raid groups. The content was accessible without being insulting or making me feel as though I was being pandered to. I didn't feel the early signs of burnout. I thought, for once, I might actually see an expansion through to its conclusion. I didn't, but not for the same reason as the previous entries in the cycle.

Elegy of the Storm

It is the middle of the night, or close enough that the difference doesn't matter. My daughter is four months old. She has been fed, changed, and is now – for reasons that remain opaque to me and, I suspect, to her – simply awake, which means I am simply awake, which means we are on the couch together with nowhere to be and nothing to do until she decides otherwise.

I have learned, in four months, to find things I can do with one hand.

The HGC tournament is on – the Blizzard-subsidized *Heroes of the Storm* premiere esports league – and it has real production value, real casting talent, and real stakes. I watch two teams draft their way through the map selection and ban phases while my daughter silently makes her way through a bottle of milk. I should be tired. Hell, *I am* tired. I am also watching the draft phase with the focus of someone who has nowhere else to be, and I think this is, in its small way, one of the better nights I've had in a while.

I had no idea, sitting there in the dark, that I was watching something that only had about two years left.

The MOBA (multiplayer online battle arena) genre did not originate with Blizzard. That is worth saying plainly, because the history of MOBAs is one of the stranger genealogies in gaming: a modification of a modification, an improvised subgenre that grew out of a custom map for *Brood War* called Aeon of Strife, which was itself adapted into *DOTA (Defense of the Ancients)* – a custom map built inside the *Warcraft III* editor. When I wanted a break from ladder matches in the 2002-2004 era, I played *DOTA*. It was a decent game. I never needed to make it my primary game, but for a match or two at a time, it held up.

(It is worth pausing here to appreciate that *Warcraft III* shipped with the tools required to effectively build an entirely new game inside of it. Blizzard gave players an engine and then moved on, and the players made something the developers never anticipated. This is what happens when a company trusts its community.)

Over the next few years, *League of Legends* and *DOTA 2* emerged as the genre's first real titans. I will be up-front about my limitations here: I never played much of either, and I am going to make assertions that are probably more impression than fact. What I can say with confidence is that both games followed the original *DOTA* formula faithfully. Five players to a team. A symmetrical map with three primary lanes, fog of war between them, jungle camps that rewarded independent play. Gold and experience accrued to whoever landed the killing blow – the last hit – on an enemy unit. Each hero leveled independently. Items were purchased from a shop. The formula was not complicated, and it worked: *League* is still going after seventeen years. The market had been established.

Which made Blizzard's absence from it all the more surprising, since they were arguably there when it was born. They had, in a meaningful sense, already built the primordial sandbox from which MOBAs emerged. But they missed the first wave entirely, or chose not to ride it, which amounts to the same thing from where the rest of us were standing. This is fine, historically speaking. First-mover advantage is real but it is not destiny, and waiting for a market to mature before entering it with a better version of the product is a legitimate strategy. Apple did not invent the smartphone, they just perfected it.

So when *Heroes of the Storm* arrived in 2015, ten years after *Warcraft III* began to fade into the background, and deep into *League*'s stranglehold on the genre, I expected it to matter. Actually, expected is probably too weak a word. I assumed Blizzard would close the gap quickly, because Blizzard had everything required to do exactly that: the deepest bench of beloved characters in gaming, the production quality of a studio that treats visual and audio craft as non-negotiable, and – most importantly – the willingness to question the formula.

Every MOBA before *Heroes* had preserved the *DOTA* architecture with modest variation. Different heroes, different skins, the same fundamental mechanical philosophy. Last hits, individual levels, purchased items, one map. Blizzard looked at this and decided, apparently without much hand-wringing, to throw most of it out.

Last hits were gone. The incentive that forced players to obsessively watch their own attack animations and count down to the killing blow – one of the genre's least fun mechanical requirements – simply did not exist in *Heroes*. Experience was shared across the team. Individual hero levels became a team resource. Gold was replaced by a talent system that asked you to choose between playstyle options at fixed level breakpoints, creating meaningful variation without requiring you to track a shop. And rather than a single map repeated across thousands of matches, the game shipped with a rotating pool of maps that played differently from each other, some of them dramatically so.

These were not small design decisions. They were a coherent philosophical rejection of the genre's received wisdom, and they made *Heroes* a more welcoming, more strategically interesting, more *fun* game to play. They also made it, as it turned out, a less successful one.

I found the game engaging in a way no other MOBA had been – and I want to be precise about why, because it matters to what comes later. The map variety was the most immediate gift: it meant that the strategic layer of the game was never exactly the same twice, that there was always something worth studying, something the draft had to account for that it couldn't account for identically every time. The shared levels and the talent system made entry approachable without removing depth. But what I particularly loved was watching the professional scene.

Watching professional *Heroes* matches is an experience that deserves to be described properly. The matches were formatted as best-of-threes or best-of-fives, and the strategy began before a single hero was deployed. The drafting phase was elaborate and strategically deep: Team A selected the map. Both teams banned a hero. Team B picked a hero, Team A picked two, Team B picked two – then another round of bans, then Team A picked two, Team B picked two, and Team A rounded out their roster. The whole structure was a hostile negotiation conducted sequentially and under pressure, like a chess clock attached to a strategy session, with both teams simultaneously reading what the other was doing and adjusting in real time.

Do you take the map your team excels at, or do you prioritize banning the map your opponents are even better at? Do you pick your preferred composition early, signaling your intentions and losing the element of surprise, or do you pick something more ambiguous and live with the flexibility that buys you? These were

the kinds of questions that defined high-level play, and the production around those questions was genuinely excellent.

Adding to this was an incredible roster of casters and commentators, both amateur and professional. Tasteless and Artosis brought over some of their chemistry from the Brood War era. Grubby, a former *Warcraft III* professional player, had deep strategic insight. And Gillyweed and Dreadnaught [10](himself a former *HotS* pro) were two of the tournament's best casters – enthusiastic, knowledgeable, with good chemistry and a real feel for the game's rhythms.

The tournament package they helmed in 2016 was the backdrop to those late nights on the couch, the thing I watched while my daughter made her infant burbles and obscure opinions known about the direction of events. I was sleep-deprived and holding a tiny person and watching professional players make decisions I found genuinely interesting, and that was enough.

In 2018, Blizzard announced the cancellation of the HGC, effectively ending the *Heroes* professional scene in one swift stroke. The development team at Blizzard was mostly redirected onto other projects. The game was not shut down, it continued to receive updates and even new heroes until 2020, but it was never the same. A competitive ecosystem, once dissolved, does not easily reconstitute, and the professional scene was the heart of the thing for me. The announcement felt abrupt. It probably was abrupt, though the business logic underneath it was legible enough.

Mike Morhaime – long rumored to be among *Heroes'* strongest internal advocates – departed Blizzard that same year. The game had not reached the viewership heights of *League of Legends* or *DOTA 2*, despite sustained investment. The calculation, made by people with more information than I had, was that continued investment was not the right use of resources. I don't even disagree. I have spent enough time thinking about resource allocation and portfolio prioritization to understand the logic.

And yet, *Heroes of the Storm* was, by any measure I care about, the best game in its genre. It was more welcoming, more strategically layered, more visually distinctive, and more mechanically creative than its competitors. It was, in many ways, a perfect encapsulation of Blizzard's long-standing design philosophy of "easy to learn, difficult to master." And on top of the design, it featured characters I already knew and cared about, maps with actual variety and strategic depth, a better spectator experience. It was, by my accounting, simply the superior product. Sadly, the better product does not always win.

This is not a novel observation. Markets do not run optimization algorithms. First-mover advantage, network effects, switching costs, inertia – all of these create conditions where an inferior product entrenches itself at exactly the moment a superior alternative arrives. *League of Legends* had a six year head-start before *Heroes* even came onto the scene, much less was trying to build a professional scene. That means enfranchised *League* players had six years of currency, heroes, cosmetic skins, and achievements already. I do not think there is a satisfying resolution to this. The game deserved better than it got.

[10] In the drafting of this manuscript, I looked them up and found that they've been recently married! I love that, that these games have the capability of bringing people together in so many ways.

I'll occasionally go back and watch some of those VODs, nearly 8 years old now. The finals of the 2017 GSC is permanently etched in my brain. Tempest vs. MVP Black, set on the Towers of Doom map[11]. Tempest went down in the early mid-game, their strategy was slow off the start and they took a while to get going, but once it did, they started to dominate every teamfight. In the end, they were able to secure all the mid-field towers, meaning shots would continue firing onto the enemy Core. Finally, MVP Black was down to one point of life left, shots continuing to fire, and we were treated to a thrilling conclusion as the last shot came out from Tempest's Core and began its inevitable journey to the enemy base. The closest analogue I can think of is watching Steph Curry heave a shot from the half-court line as the buzzer sounds, except you *know* it's going in already. And the soundtrack to it all: Dreadnaught absolutely losing his mind about what all of us were seeing.

That version of the game is still running somewhere, in the particular way that good things persist after their official ending. I'm glad I was watching when it was happening. I'm glad they made it at all.

[11] MVP Black vs. Tempest – Game 5 – Global Summer Championship; https://www.youtube.com/watch?v=u2vmZh9oKn8

You Can't Go Home Again

In hindsight, the jukebox app was what made it all work.

Not voice comms, not strategy, not loot distribution – the jukebox. A shared app where anyone could queue a song, and the raid would roll into whatever came next: something from the 80s, something from a meme, something so bad it became a special class of funny. The group had a Discord and a loose core of regulars, but it filled out weekly via the Group Finder, flexing in size to account for however many people showed up. There were no expectations, no pretension, no bench players feeling like second-class citizens. You came when you could. You were missed when you couldn't. The group would find a way. That was the version of me playing *Battle for Azeroth* in the early summer of 2019. Relaxed. Unambitious, by design.

Then *World of Warcraft Classic* was announced.

My first reaction was honest and wrong: "*Ew*. Why would anyone do that?" *BFA* was the eighth iteration of *WoW* – the original game plus seven expansions, each one refining the one before it. Streamlined talent trees. Removed friction. Better graphics, better quality of life, better everything. The idea of going back to a fifteen-year-old version of the game sounded like choosing a flip phone just to make a point. The fact that my friends wanted to play on the Pagle server[12] – named after Nat Pagle[13], the fishing quest guy – felt like the punchline confirming this was all a bit. I gave in, eventually. Not out of conviction, but as a result of peer pressure and not wanting to get left behind.

My plan was modest: level a little, see some of the old world that *Cataclysm* had bulldozed and rebuilt, and return to modern *WoW* with my priors confirmed. I chose rogue because I'd done my homework – rogue was one of the few classes whose toolkit was fully realized in Vanilla, without depending on talent rows or stat itemization that wouldn't appear until later expansions. The tier sets were actually good. Agility, one of the five baseline stats, served them well. It was the rational choice.

It was also the first class I'd ever played. Back in my friend's dorm room, years before I had my own account – an hour or two on his machine while he was in class, poking around a world I didn't fully understand yet. For this, it was the only choice that made sense.

Walking the old roads on a character without a mount – because that's what *Classic* was, forty levels of hoofing it before you even train riding, much less afford the gold cost – something shifted. I can't point to a specific moment. There's no screenshot of the realization. But somewhere in those early levels, the scales fell from my eyes, and I saw what I'd been playing for years in a different light.

This was the version of WoW I loved.

Or – to be precise – this was closer to that version than anything I'd been playing recently. And if that was true, the uncomfortable implication was that WoW had changed at some point, gradually and without my noticing, and had become something I liked less without ever giving me a clear moment to object. Similar

[12] In hindsight, I'm incredibly proud to be Aeon of Pagle. I'm incredibly lucky to call it home.

[13] It wasn't until Classic that I became aware of who Pat Nagle is, and I have to confess I'd never thought about where the Naglering item had come from until then either. One thing I've enjoyed most about *Classic* is (re)discovering things.

to how you don't notice the tide coming in until it's lapping at your feet, I'd been so caught up in it that I only noticed when my feet were wet.

Where most of my friends slowed down as the XP curve steepened, I accelerated. I was the last of my group to start, and among the first to hit 60.

What kept me moving was something that hadn't happened in a long time: the game required me to talk to people. There was no dungeon finder. If you wanted to run Scarlet Monastery or Blackrock Depths or Stratholme, you opened the Looking For Group channel and built a group the old-fashioned way – which meant talking, negotiating, reading names you'd start to recognize. It meant your reputation was a resource. You behaved like a decent person because you were one, hopefully, but also because becoming known as a jerk on a server as small as a community had real consequences.

I found a guild called Relentless Insomnia, a group that was originally from *EverQuest* progression servers, who wanted to establish an outpost here in Azeroth. It had two raid teams (each run by a different raid leader), and therefore possessed a large player base, a critical mass that mattered enormously in a game where you needed forty people to see the most important content. Aside from the convenience of having so many available people to play with, I also loved the atmosphere. With so many people, and the prospect of getting to experience the game anew, there was a palpable electricity and excitement.

Near the midpoint of phase one, well before Blackwing Lair launched, the officers approached me with an offer: the previous rogue class officer had chosen to step down, was I interested in taking over in his stead? This offer was not a complete surprise, it had been well-known for a while that I was already doing the job informally – combing through logs, tracking performance, and coaching players one-on-one. I said yes without hesitation.

It didn't last – not because I failed, but because the guild did. The guild master had made enemies around the server, which didn't help our reputation. That was a survivable problem, but then we lost our Molten Core speed ranking in the final week; we had set the server-best time the week prior, and a rival guild outperformed us in the final week, while we skipped a speedrun due to an *EverQuest* release. One week later, we came in second on BWL progression by minutes. Two back-to-back failures were at least one too many, and the internal pressure became unsustainable.

That Friday night I logged in to level an alt, and was greeted by a wall of "player has left the guild" notifications streaming down the screen. The second group's raid leader was breaking off to form a new guild, and the pitch was simple: no divided game loyalties, we'd go all-in on competing for progression and speed rankings, server-first or die trying. I called a meeting with my rogues – technically no longer my rogues, since I had no guild rank to stand on – and we talked it out. In the end, the decision was unanimous, and we all agreed to join the new team. We were the only class to move as a unit.

That new guild was eventually christened as Fully Rested, and it was, in retrospect, one of the best gaming environments I've ever been part of. We had three raid teams, each running on a separate night – not simultaneously, which was the elegant solution to a problem I'd never seen solved cleanly before: how do you give individuals the flexibility to miss a night without creating a permanent bench? FR's answer was to have no bench. You could play once a week minimum or three times if you wanted to, and every night was led by

the same raid leadership, to the same standard. If you wanted more practice, you played more nights. The only cost was your time.

We took performance seriously. I'd never been a part of a team where everyone was as motivated to perform as I was, and it was a liberating experience. What I found particularly pleasant about it was that there was never a question about what we were about, the expectations were made clear when people joined, and for the most part, people lived up to them, or self-selected to retire or leave the team, but there was never any drama about it. And with a team as aligned as we were, it came as no surprise that we reclaimed our speed ranking crown before the end of BWL. And for the final two phases – Temple of Ahn'Qiraj and Naxxramas – we claimed the top speed ranking and progression ranking for Pagle, and in the case of the latter, we came in fourteenth in the world.

Around the same time, I was convinced to level and play a mirror alt: a second character, same class, same spec, designed not for variety but for volume. In a game where your raid lockout resets weekly, your opportunity to practice on any given encounter is capped. A second rogue meant a second lockout, a second set of raid groups, twice the repetitions. We were vendoring rogue loot by then, so gearing a second character happened quickly. That extra practice paid dividends, and the growth I experienced over those months – with twice as many opportunities to test and gather data leading to refined positioning and optimized cooldown timing – was some of the steepest improvement I'd managed in any game, at any level.

They say you can't go home again. Thomas Wolfe meant it the hard way: home isn't a place, it's a constellation of conditions – people, time, your own younger mind – that can never be perfectly reconstructed. The river has moved on. The water you stepped in is already gone.

That's true, but the funny thing about *Classic* is that I didn't go there to go home.

Remember, I'd gone there initially as a joke, a meme. It was something I'd done to say I'd done so. And what I found, walking those old roads at walking speed, was not nostalgia – not a longing for who I was at twenty-one playing a rogue on my friend's computer while he was off at class. What I found was a mirror that showed me the gap between the player I'd been and the player I'd always wanted to be, to see that gap and get an opportunity to walk old roads with the wisdom and perspective gained through age and experience. It was if the game was daring me to quit talking about things I'd wished I'd done, and to actually go out and show everyone (but primarily myself) what I can actually do.

In Vanilla *WoW*, I was a nobody. I never killed Ragnaros. Never set foot in Blackwing Lair. Never saw Naxxramas until a much later expansion repurposed it for lower-level players. I was somewhere in the broad, undifferentiated middle of a massive game, playing casually, never quite finding the edge of what I could do.

So while I didn't go back to *Classic* to relive the glorious days of my youth (because they weren't glorious), I could go back with the addition of age, wisdom, and perspective. *Classic* became an experiment, where I could seek an answer to the question: "What could I have done back then, if I was the player I am now?" And every week, I ask that question anew.

You can't go home again … but you can absolutely go back to Azeroth.

The Burden

The call to leadership often comes before you're ready for it.

I had spent *Classic TBC* as a passenger – deliberately, happily so. I'd served as an officer in Fully Rested, carried a portion of the weight, been useful in the specific ways that role officers are useful. I had parted ways from FR for *TBC*, because I wanted to be able to focus entirely on my play; I wanted to compete, wanted to master my class without the overhead of managing other people. I did all of that, and I don't regret a single minute of it. But somewhere in the long tail of the Sunwell, while the expansion went down the final stretch of the inevitable march toward *Classic Wrath*, I could feel something tugging.

My wife Clover felt it too. One evening in the pre-patch interregnum before *Wrath* arrived, we talked it through properly and arrived at the obviously crazy conclusion: we were going to build a raid team from scratch, in the four weeks remaining prior to the launch of *Wrath*, and we'd do our damnedest to make it competitive with the best teams on the server.

This was, by any rational measure, a lunatic proposition. A brand-new team, competing against established guilds with rosters and histories and two full expansions and years of social infrastructure, with only a month to do so. But we had things going for us that we perhaps undersold. The two of us had two expansions worth of reputation as top-tier players. We had a network of connections on Pagle wide and deep enough to generate a roster if we worked at it. And we had each other – a genuine partnership, not just two players who happened to be in the same guild.

A friend reached out with a lead on the infrastructure piece. On Pagle, a guild called Triumvirate had a history of hosting multiple independent raid teams under one banner; they'd had three at launch in *Classic*, but were down to one heading into *Wrath*. They had room, and they had structure, and while we'd be our own raid team, the Triumvirate name would confer a certain level of legitimacy that a new group would need. After a single conversation with the GM, it was obvious what the answer was.

We still needed a name for our team. Clover suggested Starcaller. Nailed it in one.

Our responsibilities diverged from the start, cleanly and naturally. She took the lead on the human side of the operation: recruitment, public relations, the endless texture of social management that keeps a raid team from dissolving into a collection of strangers with a shared calendar invite. I took the strategic side: in-raid tactical leadership, long-term planning and goal-setting, performance analysis, composition planning. I would play a prot paladin as our main tank. She would play retribution paladin and serve as melee lead. Neither of us had done this exact job before, but we'd find a way to make it work.

The pre-patch ran four weeks, and they were filled with activity. Recruiting. Planning. And for me, hosting pickup raids in a nerfed version of Black Temple two or three nights a week – not because the content mattered anymore, but because *I* needed the practice raid leading. There is a meaningful difference between issuing the occasional call from the middle of a raid and speaking from the bridge, watching everything, directing everyone, being the voice that twenty-four other people are listening for while simultaneously doing your own job. I had never held that much responsibility at once. I needed to learn how to carry it before the real content hit.

Those raids paid unexpected dividends. The other Triumvirate team, already established, sent members to our pickups; we went to theirs. Camaraderie formed without anyone deciding it should. By the time

Naxxramas went live, we weren't just two teams sharing a guild tag – we were neighbors who'd already started borrowing things.

And then the miracle happened: everything worked out!

By launch night we had a full team. Good players, with a decent raid composition. We'd had some opportunities to play together, and so we wouldn't be starting from scratch with Naxx. We'd self-organized into dungeon groups for leveling and gearing in preparation of the first raid. And the stars aligned in the way they almost never do, and Naxxramas fell cleanly, without the catastrophic first-week collapse that ends so many new teams before they can find their footing. To keep our team from getting complacent, we spent the back half of Phase 1 not just clearing content but *improving* – running speedrun weeks, pushing our limits for the pleasure of seeing exactly where our ceiling was.

There is an achievement in Naxxramas called The Immortal. In a single raid clear, the group raid must clear every boss without a single player death. No errors, no bad luck, no lapses. There are many teams that never get it.

We earned it twice.

I was adamant about the need to complete it a second time. Like most 25-man teams, we kept a small bench – every week two or three players would rotate onto the bench, to provide a buffer for absences and last-minute callouts. I argued that everyone was part of the team, we all had put in the same work, and I wasn't interested in a version of success that left any of them standing outside it. For the second clear of The Immortal, everyone on the roster who wanted it, got it. That meant more to me, honestly, than getting it the first time. Once can be a fluke, but twice is a trend.

We also earned a 15th-server speed ranking for Naxxramas – which sounds modest until you account for the fact that the majority of our players had never attempted a speedrun of anything before joining Starcaller.

Then Ulduar arrived, and the expansion showed its actual face.

The hardmodes changed everything. Blizzard's design philosophy – easy to enter, difficult to master – crystallized in Ulduar in a way that Naxxramas had deliberately avoided. We had planned to enter the first week and clear four or five hard modes on the strength of preparation and grit. We cleared zero.

I do not think I could have scripted a better test of what we'd built, and our players passed it. It would have been easy for disappointment to get the better of us, for players to quietly seek greener pastures. Instead, our raiders chose to dig in, to learn, to get better, and to trust the process. We found our stride. We finished Ulduar with every hard mode completed – every single one.

I should have been prouder of that than I was. What I was instead was tired.

This is where I have to be honest about the thing I failed to do, the absence that I could not fill and worried about constantly: I never found a suitable lieutenant.

A raid leader who cannot be replaced is a single point of failure with twenty-four other people depending on it. I knew this from the beginning. I tried throughout Phase 1 and into Ulduar to identify someone – anyone – who could step into the tactical role if something happened to me. Nobody with the willingness and the capability appeared together in the same person. I found excellent role leads: a strong caster lead, and a healer lead I trusted (and depended on) completely. But the captain's chair was mine alone.

It's lonely at the top.

In 25-man raiding that's already a significant burden. In Ulduar it became something else, because the 10-man version of the same content was functionally required – not optional, but necessary for class-specific gear, for trinkets, for the pieces that rounded out a character's power. Twenty-five raiders means you need to run multiple 10-man groups each week. And they needed to be led. And I felt the obligation.

So I led them. I had to.

Trial of the Crusader was simpler content – shorter, more focused, lower difficulty ceiling. That should have felt like relief. Mostly it felt like a chance to catch my breath before Icecrown.

And Trial gave us the second great group achievement of the expansion: A Tribute to Immortality. Like The Immortal in Naxxramas, it requires a full clear without a single death. By contrast, it requires heroic difficulty – whereas Naxxramas only had normal difficulty – which elevates raid mechanics to the point where they will simply kill a player if they're not managed correctly. The difficulty was an order of magnitude higher.

Again, there are many teams that never complete this. But as with The Immortal, we completed it twice, and everyone on the roster participated. I refused to count it otherwise.

The final full raid dungeon, Icecrown Citadel, is more like a prolonged siege than a quick battle. It features twelve bosses strung across a fortress of undead, floor after floor of material to learn and re-learn, it was not the longest raid in *Wrath,* but it was the most difficult of any raid tier to date. It was not easy, but we got there in the end, and we defeated Arthas on heroic difficulty before *Wrath's* sunset – a goal I had held privately from the first week of Naxxramas, never said aloud in case I had to eat the words.

There is technically one more raid in *Wrath*: Ruby Sanctum, a single-boss dungeon that arrived in the final weeks. Starcaller did not achieve the heroic kill. By that point the expansion had simply run out of runway, and so had we. I think we could have gotten there, but our raiders were tired, and the will simply wasn't there anymore. We cleared it on normal, declared a moral victory, and moved on. I am at peace with it.

I learned more in the year and a half of leading Starcaller than in any equivalent stretch of my professional life.

I learned what it costs to sit in the captain's chair with no relief crew. There was a week somewhere in the middle of the expansion – I can no longer pin the exact timing – where I had COVID. Not badly, but badly enough. I raided anyway. I didn't feel like I had a choice. The fight could not happen without me, and I could not explain why not in a way that didn't reveal the vulnerability I'd spent months trying not to have.

There was nobody else.

I learned about the roster boss – a problem that has no mechanical solution, only vigilance and luck. Even with a healthy bench, you are always three or four departures from non-viability. When a key player leaves, they sometimes take a friend. When two players leave simultaneously, it can destabilize the whole social structure of the team. I spent a year playing a game that was not entirely *Wrath of the Lich King*. I was also always playing *keep the roster intact*.

I learned what it means to hold tactical responsibility for twenty-five people at once. An individual player tracks themselves. A role lead tracks their four or six players plus the broader fight. A raid leader tracks all

twenty-five, plus themselves, plus the fight, plus the performance of every role, plus the calls that need to go out in the next three seconds, plus the strategic adjustment that might need to happen after this pull. The cognitive surface area is vast, and it is yours alone, and you have maybe a second to process any given piece of it.

And yet, had I known all this going in, I still would have done it.

Here is what I know about that year, looking back with the clarity that distance affords: I'd estimate that ninety percent of our raiders performed beyond what their natural skill level, taken alone, would have predicted. We did this by being more than just a raid team, we built a community with conditions in which people discovered something better in themselves than they had arrived with[14]. I have watched players clear content they told me, at recruitment, they weren't sure they were capable of. I have watched people who came to us as role-fillers leave as officers.

Clover and I announced our retirement from leadership before the expansion formally ended. We weren't secretive about the reason: we were done, fully, and we wanted to be able to direct our energy back into playing. We made it clear we'd stay as individual contributors if anyone wanted to continue under a new leader. Nobody stepped forward. Starcaller retired as an active raiding team and became a community of friends, which is a better ending than most teams get.

In the end, Atlas shrugged. I had not realized, until I set it down, how heavy that burden had actually been.

[14] One of my favorite leadership aphorisms comes from Simon Sinek, who says, "The real job of a leader is not to be in charge, it's to *take care of those in our charge*." I hear the phrase 'servant leadership,' but I think the real trick is 'supportive leadership.'

Flow State

I opened WarcraftLogs this morning and checked my ranking.

Ninth place.

Yesterday it was eighth. Last week, fifteenth. The week before that, twenty-fourth. I've been watching the number climb the way you watch a stock price – obsessively, superstitiously, knowing that attention alone doesn't move it but checking anyway, because the alternative is not knowing.

Some context is useful here. We're talking about retribution paladins in Heroic Throne of Thunder – *Mists of Pandaria Classic*, the current tier. Twenty thousand, five hundred and thirty-six ret paladins have killed at least one boss in this raid. At 9th, that means I am in the top percentile (technically, the 99.9th percentile). Filter it down to North American players: fourth. Alliance players: third. But the global number is what pulls at me, and the global number is ninth.

For now.

I've been circling this altitude for a while without ever fully landing on the summit. I peaked at second place briefly during Black Temple on the rogue – before getting pushed back to nineteenth by the time the logs partition closed. I played a tank in *Wrath* – and while I started the expansion playing more defensively-minded, as my comfort with my raid team grew and I could see where I could push the envelope, I ended up in the mid-hundreds among tanks in Icecrown Citadel. In *Cataclysm* I swapped to hunter for Firelands and finished world forty-sixth, then took the rogue again for Dragon Soul and landed at forty-fourth. And in *Mists*, back on the ret paladin, Heart of Fear and Terrace of Endless Spring: fortieth.

It's been a rollercoaster, frankly. It feels like each expansion, and each phase, I'm constantly shifting, always adapting. Even so, the current state of being balanced on the cutting edge, feels like something different altogether.

Part of it is structural. In *Wrath*, I was leading the raid team – tracking cooldowns, calling positioning, managing thirty people's energy levels and interpersonal conflicts alongside my own performance. There's a tax that comes with that responsibility[15]. It isn't just the mental bandwidth consumed by logistics; it's the perceptual overhead of holding the entire team in your head at once. You can't be inside your own play when you're also responsible for everyone else's. The two modes of attention don't coexist easily[16].

What I didn't expect was what happened when I stepped back from leadership.

When I returned to individual contributor in *Cata*, I didn't just shrug off the burden of leadership – I kept the awareness and perspective. A raid leader has to see the fight at a level your average player never develops. They track the whole board: where the boss is moving, which cooldowns are available across the team, what's about to happen and what needs to happen in response. That spatial, temporal comprehension doesn't disappear when you take the headset off. It compounds.

---

[15] I've been talking about this for years as the "raid leader tax," a phenomenon where most raid leaders I know can't put out their best individual performance in groups they then have to lead.

[16] This is actually a good analogue for program management, illustrating the difficulty of keeping a high-level 50k foot view as well as a low-level task-oriented perspective in mind at the same time. It's hard.

Think of it like training for a sprint with weights on your ankles. The weights don't make you a better sprinter by making you faster. They make you better by making you *stronger* – so that when you remove them and run unencumbered, you move in a way you simply couldn't before.

I had spent an entire expansion building that muscle. In *Cata* and beyond, I've been able to finally use it entirely for myself.

There is a hierarchy to this game that most players never articulate.

At the bottom, there are players who don't fully understand how to play their own class – the basic rotation, the core mechanics that make them effective. Above them: players who've mastered their class but get confused when the fight itself demands their attention. Above them: players that hold both in mind simultaneously (these are good players). And highest, the great players know which rules are absolute, which can be bent, and which can be safely broken, and know which is appropriate at any given moment.

But the truth is, truly world-class players don't think about any of this in the heat of the moment, at least not consciously.

Not because they've stopped caring, but because they've internalized everything to the point where it runs beneath conscious attention. The rotation is automatic. The fight mechanics have become reflexive. The whole complex apparatus of decision-making that less experienced players are consciously managing – they've pushed it down below the surface, and what emerges on top is something like a kind of clarity.

Let me give you a sense of perspective. I play ret paladin, which means I have six primary combat abilities, three primary damage cooldowns, and a host of utility spells. My global cooldown is approximately one second, which means I need to be prepared to press a button every second to play optimally. Add in internet latency and the physical neural latency from stimulus to response, and we're talking literal split-second thinking. What's more, I find that in the span of a single GCD, I can feel my eyes rapidly shifting from tracking ability CDs, boss abilities to dodge, timers tracking upcoming events, often cycling three or four times in that time.

When I'm deep in a raid and performing at the edge of my ability, there's a state that comes over the play. I know what's happening right now. I know what's *about* to happen. I'm already sequencing abilities in my head several GCDs into the future. The gap between intention and execution collapses. Time feels different. This is what people mean when they say someone is "in the zone"– it's not the absence of effort but the disappearance of friction. You're doing the thing without being aware of doing the thing, and yet you're somehow more present in it than you've ever been.

Flow state.

Maybe it's the class. The ret paladin has been my primary identity in this game for over a decade – the Holy Power system, the judgment calls on when to spend and when to hold, the burst windows that reward patience and prior planning, and punish waste. It fits the way a well-worn tool fits a hand. Maybe it's repetition: prior to *Wrath Classic*, I played multiple main classes in different raid groups; in *Cata* and *Mists,* I've changed classes between phases but played mirror alts in my raid groups.

Or maybe it's that I finally understand what it actually takes to be great.

Here is something counterintuitive about world-class performance in a game people still call a hobby: you have to study it. I spend real time between raids in the logs – my own and other people's – the same way a basketball team watches game footage. I'm not looking for big mistakes. I'm looking for the ten decisions per minute that collectively produced a result, trying to understand which of them were necessary, which were optimal, and which were just what I do habitually because I've never interrogated them.

And just like a sports team: I can't do that work alone. *WoW* is not a solo endeavor. It looks like one on a chart – one name, one number, one ranking – but that number is produced by twenty-five people doing everything right at the same time. I have an almost identically equipped paladin playing in a separate raid group, same pilot, and that character sits at 182nd on the global list. The gap between ninth and 182nd is not my skill. It's context. It's team. It's what the raid around me enables me to do.

Three of the top five North American paladins play in my guild.

I find that simultaneously hard to believe and yet impossible to explain any other way. We compete intensely with each other – the kind of internal competition that sharpens rather than fractures, because we also talk. Every week between raids, we compare notes, share ideas, test adjustments and report back. The unique advantage we have is that we can run true apples-to-apples comparisons: same strategies, same fights, same duration, different executions. We can locate the variance and understand it. That feedback loop is something most players don't have access to.

The lesson I keep arriving at, from every angle: if you want to perform at the highest individual level, you have to invest in the collective intelligence of your team. Not because you're responsible for everyone else, but because understanding what the whole system is doing is the only way to fully understand what *you* are doing inside it. You probably see things the leadership doesn't. Share what you see.

This is, I realize, not only a philosophy about raiding.

Ninth place today.

I checked again this morning, and it's still ninth. Maybe tomorrow it will be tenth. The number moves like a living thing – other people continue to compete and improve, doing their own analysis, getting better in ways I can't observe until after they've already happened. The ranking is a snapshot of a competition that never stops running. Frankly, I don't want it to stop.

Something happened to me across these years – across the raid leader's burden, the class switches, the long study of other people's play – that I didn't plan and couldn't have engineered directly. I became good at something difficult. Even if you account for the fact that I'm looking at one specialization, of one class, in a single raid tier, of a specific incarnation of a game that is now 15 years old: *it's rare to be able to say, "I'm the ninth best X in the world at what I do," and be completely accurate about it.*

Not as a side effect of enjoying it. Good because I treated it like it deserved to be taken seriously. Because I believed that within whatever small domain this is, excellence was worth the effort it required.

The number on the screen is just evidence.

Remembrance of BlizzCon Past

The morning light comes in low and honest, the way it only does in the hour after school drop-off, when the house has exhaled and the day hasn't yet made any demands. I sit down on the couch, just for a minute, just to gather myself before the day's work begins.

My eyes drift, the way eyes do when a mind is briefly without instruction, and land on the coffee table. There's a stack of books there. A vintage illustrated copy of *The Hobbit*, a copy of *Harry Potter and the Sorcerer's Stone* of similar type and provenance. And on top, as it has been since the day we moved into this house, a large-format art book from a collector's edition of a video game expansion released in 2007.

By conventional standards, it shouldn't be on the coffee table. Most people in their middle years have coffee table books about their interests, but they're typically about more mainstream topics: rock music, impressionist art, vintage cars. And yet I never considered placing this particular book anywhere else in my home.

I reach forward and lift the cover. The inside page is scattered with signatures – scribbles of silver paint pen, dozens of them, maybe more, scrawled across the art and margins and white space. Names I know. Names I chased down across two days of a convention I'd flown halfway across the country to attend, sleep-deprived and twenty-two years old and completely certain this was a reasonable thing to do.

The morning light catches the silver, and the signatures glitter.

I lose myself in the pages for a little while.

My friend Anton came to me in the spring of 2007 with a question that was really a proposition: had I heard of BlizzCon? I had not. The first one had happened in 2005, and I'd missed it entirely – not through negligence, but through the simple arithmetic of being young and broke in Texas, a significant distance from Irvine, California, where Blizzard Entertainment makes its home. Even if I had known about it, I couldn't have gone. I didn't have money for a *World of Warcraft* subscription back then, let alone a plane ticket.

Anton made his case anyway. Here is what he proposed: a red-eye flight into LAX, cheap, landing eight hours before the convention opened. Drive to Anaheim. Sleep in the car. Two days of convention. Then, continue the trip by crashing on the couch of his brother, a student at Caltech in Pasadena. In short, a week's vacation that would cost roughly the price of a late-night plane ticket, a small economy rental car, one unfashionable motel room, and a $125 convention ticket.

I was twenty-two years old. I said yes immediately.

We landed a little after midnight. Collected our bags, caught the bus to the rental center, picked up our car, and made the only logical first stop: a late-night pilgrimage to In-N-Out Burger (a delicacy that had not yet reached Texas and therefore carried the full weight of myth and rumor). It was probably well after 2am by the time we got back in the car and drove south on the I-5 toward Anaheim, found distant parking, reclined our seats as far as the Hertz econobox would permit, and attempted to sleep.

I am quite certain that attempting this today would be a disaster. My body would file a formal complaint. My brainstem would mutiny. At twenty-two, it was merely uncomfortable, which is a category of experience the young navigate without much trouble.

At 8am we got in line at the convention center.

I remember looking around at a roiling sea of people and feeling something that I hadn't anticipated: recognition. I didn't know anyone there except Anton, and yet I immediately felt as if I knew these people. The particular body language of someone who has strong opinions about which *Warcraft III* race is most imbalanced, or who their favorite *WoW* raid encounter was. There were people dressed up in costume, and even the regular people wore t-shirts advertising games and guilds and inside references legible only to a specific kind of person. I was among my tribe, and I knew it before a single word was spoken.

I had no idea what to expect from the convention itself. Less information existed about everything in 2007, and this was only the second BlizzCon ever held – they'd skipped 2006 entirely. So when I checked in and was handed a bag of things, I was genuinely surprised. Inside: pins, magnets, a pack of trading cards. A shirt. Coasters. A card with a code for an in-game toy – a Murloc costume, which was exactly as absurd and delightful as it sounds. And another card, for something else entirely: access to a beta for an upcoming Blizzard game. Jackpot.

Also in the bag was a bookmark with a quest on it. Find a location on the convention floor. Talk to the NPC. Get a stamp. Complete the quest, collect a reward. It was a small thing, a bit of manufactured purpose – and it worked completely. I am apparently still susceptible to quest design in the physical world.

This was the convention where Blizzard announced *Wrath of the Lich King*. Arthas, the Lich King, one of the most beloved villains in the entire Warcraft mythology, would finally be given his due as an endgame confrontation. It would be the logical conclusion to a story arc spanning five years, two games, and three expansion packs. And my primary game at that point in time had become *WoW,* so by any reasonable measure, this should have been the event of the convention for me.

Not a chance. The playable demo for *StarCraft II* was on the floor.

I was vaguely aware they'd announced the game a few months prior, at an event in Korea – very appropriate for that series. A sequel to a beloved franchise that sat at the absolute pinnacle of the real-time strategy genre was not a surprising development, in retrospect. But knowing something is coming and actually sitting down to *play* it are different experiences entirely, separated by a gap that no amount of anticipation can bridge.

Anton and I fell into a loop that consumed most of the first day: wait in line for a demo station, play until the machine kicked us off, exit the station, immediately rejoin the line. We did this without discussion, without ever seriously considering doing anything else. The game was electric. We were moths and it was a very bright light.

When we'd finally burned through enough of the *StarCraft II* demo to face the rest of the convention without withdrawal, I turned my attention to the Meet the Devs tables.

I had brought my *Burning Crusade* Collector's Edition art book with me from Houston. This was not an accident. I had a plan: periodically cycle through the signing tables, find every artist and developer I could, and get them to sign it. If they hadn't worked on *Warcraft*, I offered them the BlizzCon program instead. The signatures were almost beside the point. What I was actually doing – though I wouldn't have articulated it this way at twenty-two – was making human contact with the people who had made the things I loved. Reaching across the distance between audience and author, consumer and creator, and briefly collapsing it.

I kept that book through every subsequent chapter of my life. Years later, when I met the woman who would become my wife, I eventually gave it to her – or rather, gave it to us. It lives on our coffee table now.

The Blizzard Invitational was my first live esports experience, and it did not disappoint.

I'd watched professionals play *StarCraft* and *Warcraft III* before, through casts and replays, the way you watch anything when you can't be in the room. Watching it live is a different thing altogether. The crowd reads the game in real time. The tension moves through the room like wildfire. When Grubby – one of the most recognizable personalities in the *Warcraft III* scene, a player I'd followed through replays and casts – lost in the semifinal, the collective deflation was palpable. A room of strangers, briefly united in disappointment.

I hadn't been following the *StarCraft: Brood War* pro scene closely in 2007. Several years of deep immersion in Warcraft had pulled my attention in other directions. But even a casual familiarity with Korean professional *StarCraft* was enough to understand what it meant to watch sAviOr play in person. He was the only Zerg *bonjwa*[17] – a player of such sustained, dominant excellence that the word "champion" is insufficient. Watching him dismantle Nal_rA in two straight games did as much to reignite my love for *StarCraft* as the demo I'd been playing all day. I was marginally disappointed at how quickly the series ended, but mostly I was just grateful to have been in the room when it happened.

The concert that night was Video Games Live – orchestrally arranged video game music, which I already loved – and Level 70 Elite Tauren Chieftain[18], Blizzard's in-house metal band, which is exactly what it sounds like and exactly as wonderful. Jay Mohr was the master of ceremonies, a comedian I knew by name but not by work. His opening material bombed. The crowd, which had plenty of opinions and no particular incentive to be polite about them, let him know. And then something unexpected happened: he stopped trying to win them over and started giving it back. The heckling became a conversation. The conversation became something genuinely funny. The room, which had been resistant, came around completely.

I've thought about that moment more than once since then. There's something instructive in it – about reading a room, about abandoning a plan that isn't working, about finding the actual audience in front of you instead of performing for the one you imagined. A comedian who lost the crowd and won it back by deciding to be honest about losing it in the first place.

Anton and I stayed in California for at least another week after the convention, visiting his brother, doing things I can no longer remember. The trip continued. Life continued. And yet when I try to retrieve that week, there is almost nothing there – just a pleasant blur of couch-sleeping and Pasadena afternoons, thoroughly overshadowed by two days of convention.

Less than two days. Sleep-deprived, slightly broken-backed from that nap in the car, twenty-two years old and running on caffeine and adrenaline and the particular energy of being somewhere you felt like you were always supposed to be.

---

[17] This is a fun, obscure word, commonly used to describe a player with a sustained level of dominance. There are no defined metrics, it's more a consensus view from the community of who is "the best" of their era.

[18] Apropos of nothing, I love that this has become an actual character in *Heroes of the Storm*.

There is a particular hour in the afternoon when the sun finds my office.

It varies by season – the orbital geometry of our planetary system shifts, the angle changes, and the light arrives a little earlier or later depending on where we are in the year. But it comes, and when it does, it catches a specific pane of glass, the frame protecting a painting that is hung above my desk. I look up, and there he is.

Zeratul.[19]

Arm extended, psionic blade ignited, cloak trailing behind him in a breeze that exists only in the paint. He is small in the frame – deliberately, meaningfully small – dwarfed by the scale of the Xel'Naga temple he is moving through. Ancient architecture looms around him, the ruins of the race that created his own, and he is alone in it, blade drawn, pressing forward into the unknown. I have looked at this painting hundreds, maybe thousands of times. I still find myself asking the same question.

What is he feeling in this moment?

Awe, almost certainly. He is walking through the remnants of gods. Wonder is the only sane response to that. But fear too, I think. He knows what else might be in those ruins. He knows that the Zerg are never far, that Kerrigan – the Queen of Blades, the most terrifying creation the Swarm has ever produced – could be waiting anywhere in the dark. In the cinematic that this painting echoes, she is. He rounds a corner and there she is, menacing and dominant in the way that only a villain assured of their invincibility can be. Any reasonable calculation of the situation suggests that Zeratul is doomed: he should not be here, he should not have come, he cannot hope to win this. He raises his blade and leaps anyway[20]. There is a type of courage that doesn't require certainty. That looks at an impossible situation and decides, without much deliberation, that taking the leap is still the right thing to do.

The afternoon light shifts. I continue with my work.

The first time you do something extraordinary, everything is new to you; the possibility space is endless and wide open, leaving it up to you to discover what is and what can be. The second time you do something extraordinary, you know what to expect, because you've been there before. That knowledge is both a gift and a thief.

I skipped BlizzCon 2008 for the same reasons I'd nearly skipped 2007, the harsh arithmetic of disposable income failing to support ambition. I returned to BlizzCon 2009 because I had a new job, a little more money, and a determination to recapture something I hadn't fully understood the first time. I was older by two years, which in retrospect is not very much time at all. Old enough to arrive with intention rather than just wonder.

The convention had grown. Four full halls now, twice the footprint of 2007, and yet it had the same density and quality – everywhere you looked, there was something worth looking at. *Cataclysm* was announced, the next chapter of a *Warcraft* story I'd been living inside for the better part of a decade. *Diablo III* was playable on the floor. These were not small things, and yet they settled into memory the way the

[19] I've been trying to find the name of this painting and the artist for a long time now with no success, it drives me nuts.

[20] A favorable comparison to Han Solo in *Empire Strikes Back* is the scene when they are surprised by Darth Vader, and without thinking he draws his blaster and starts firing.

second viewing of a great film does - recognized, appreciated, but ultimately not revelatory. The spark had become a flame. Steady. Warm. Reliable.

What I remember most about 2009 is not anything particular that happened at the convention. It was the comfortable embrace one feels upon returning home, to be amongst one's people again.

In 2010 I came back again, and this time I brought home two things that have never left.

The first is a statuette of Deathwing – Neltharion the Earth-Warder, the corrupted black dragon aspect, rendered in miniature and included in that year's swag bag. He occupies a place of honor on my bookshelves now, front and center among the Lego sets and gaming collectibles and other artifacts of a life spent caring about fictional worlds. He has survived every move, every reorganization, every periodic reckoning with what deserves to stay and what doesn't.

The second thing I brought home from BlizzCon 2010 is steeped in much deeper meaning.

The Child's Play charity auction occupied a small piece of the convention floor, and I drifted through it the way you drift through a museum – not looking for anything specific, just interested in the works of art, and open to being stopped by something.

That's where I found it. A framed digital painting of Zeratul, clearly a work of concept art for *StarCraft II*. I stood in front of it and knew, with the particular certainty that bypasses logic and deliberation entirely, that I had to have it.

I told myself it was for a good cause. Child's Play does genuinely good work – bringing games to children in hospitals, to kids for whom play is not a luxury but a lifeline. But if you demanded complete honesty, I would tell you: the charitable reasoning is *ex post facto* rationalization. The truth of it is, I saw it and I had to have it.

I was young. I was making entry-level salary with entry-level savings to match. The justification was simply a way to rationalize spending a significant percentage of my income on something so extravagantly frivolous. I bid anyway.

The final hour and a half of the auction I spent in a particular kind of focused anxiety, drifting through the other pieces, with one eye always back on mine, running contingency calculations in the event that someone else – some heartless monster with deeper pockets and no regard for my need to have this artifact – decided to outbid me. Someone did, repeatedly, and with what I can only describe as a cheerful disregard for my emotional state.

The last ten minutes were tense, in the way that only auctions and close sporting events manage to be tense – a specific, compressed, very personal variety of suspense. I had to thread the needle between responding too quickly, which signals desperation and invites escalation, and waiting too long, which risks losing the piece entirely.

In the end, I won. I handed over my credit card information. Gave them my address for shipping. The prize was mine.

The painting hangs above my monitor now. Has for years. It is the first thing I see when I look up from work, and the last thing in my eyeline before I close the computer for the night. I did not know, standing in that auction corner in 2010, that I was buying something I would carry into every subsequent version of my life. I did not know what was coming – the years of drift, the marriage that would slowly extinguish things I

cared about, the long dim interval between the man I was at BlizzCon 2010 and the man I would eventually need to become again.

At the time, I just knew I had to have it.

There is something to be said for the instincts that operate below the level of reasoning. They are keeping score even when you aren't paying attention.

It is January 2021. Ten years have passed since BlizzCon 2010. In that decade I accumulated the artifacts of a shared life with someone who turned out not to be my person — furniture and kitchenware and the accumulated objects of a household that was never quite a home.

My daughter, now four years old, is visiting and playing with my parents. My soon-to-be ex-wife is absent for a different reason: our marriage is in the throes of a prolonged death spasm. I am alone in a house that is quiet in a way that borders on oppressive, the same quiet you'd find in a crypt, or an abandoned museum.

I am packing boxes.

The physical part of packing is straightforward enough. You find boxes, you fill them, you seal them, you stack them. The real skill lies in the act of deciding. Every object you pick up asks the same question: does this come with me? And underneath that surface question, the real inquiry, which you do not ask aloud because there is no one to ask: who am I, on the other side of this? What does that person need? What does he carry forward, and what does he set down here, in this room, and walk away from?

I find, when it comes to it, that I do not need most of it.

What I do want fits in surprisingly few boxes. Clothes. The necessities. And then the things that belong to an earlier version of me — a version that had not yet been talked out of his enthusiasms, had not yet allowed the things that lit him up to quietly dim under the weight of a life that had drifted badly off course. A book, heavy with signatures. A statuette of a dragon. A painting, wrapped carefully and set aside before I'd packed anything else, the only piece of art I take.

I was twenty-two when I first understood what it felt like to be exactly where I was supposed to be. I am thirty-six now, and I am trying to remember what that felt like.

The boxes are a question. The answer, slowly becoming clear as the afternoon wears on and the rooms empty, is simpler than I expected.

You carry forward what still has meaning. Everything else is just weight.

I met the woman who would become my first wife just a few months after BlizzCon 2010. I was in my middle twenties, not anxious to settle down but aware, in the vague way you become aware of things at that age, that it was probably time. We met, we dated, we got to know each other. She made token overtures about liking *Final Fantasy* on the Super Nintendo, seemed to enjoy some of the same fantasy and science fiction I did, and I figured she was one of the tribe.

I was wrong, but I wouldn't understand how wrong for years.

When you are new in a relationship, it's easy to mask off portions of yourself. She hid her distaste for video games. I hid how important they were to me. We met somewhere in the middle, which sounds like a reasonable compromise until you understand that in a long-term relationship, that middle ground is where things go to slowly die.

The truth is she didn't mind games categorically. Single-player games were fine – something you engaged with on your own time, paused when needed, closed when called upon. Clean. Contained. No fuss. The friction started the first time I booted up something multiplayer, something with other people in it who were depending on me to be there. She couldn't understand why I couldn't simply stop. To her, the other people in the group weren't quite real.

It was a friction point that might have been managed. And then *Mists of Pandaria* released. I remember the moment she saw what I was playing. The grimace. The clenched teeth. The words: *"what is THAT you're playing?"* For whatever reason, she had an intense personal dislike of the *WoW* franchise. It was reflexive, irrational, and absolute.

Though her feelings would soften somewhat over the years, I became skilled at tiptoeing – finding the edges of what was acceptable to her, keeping the games I loved at a careful distance from the parts of my life that required peace. I could play the games I loved, but that always came at a cost. I could never engage to the extent I wanted, and never enough to recapture that sense of real mastery, strategic or mechanical. The flame didn't go out. But it was banked low, starved of the oxygen it needed, kept alive through careful and diminishing tending.

Ten years is a long time to tiptoe. In the end, it was simply too much.

I met Clover in the late summer of 2020.

She was a very talented Paladin healer, a new recruit to Fully Rested on the Pagle server, who joined us partway through the Temple of Ahn'Qiraj while the rest of the world was doing … whatever the rest of the world does in the grips of a global pandemic. I was an established senior officer, vaguely aware of her name from the PUG/GDKP communities on Pagle, but I didn't know her personally.

She first knew me as Aeon – the name I'd been carrying in these games for years, the name under which I write. We knew each other's in-game personae before we knew the person behind the keyboard at all, which turns out to be a more honest introduction than most.

Over the following months we got to know each other properly. We were of an age, our childhoods running in chronological parallel – the same games at roughly the same times, the same cultural landmarks, the foundations of personal identity. We shared a love of cooking and good food. And, crucially, a passion for *World of Warcraft* and a burning desire to be among the best to play it. We went from guildmates to friends in short order.

And we were both, as it happened, in unsatisfying marriages and contemplating futures without our spouses, though we didn't discuss that directly until much later. What we had first was simpler: a friendship, and a shared language, and the particular ease that comes from talking to someone who already understands you, even before you say anything.

Most normal people, wanting to express something warmer than friendship, would gift someone flowers or chocolates. I did what any reasonable person would do: As we battled our way through the Qiraji armies, I would flirtatiously master-loot [Huge Venom Sacs] directly to her inventory.

She thought it was funny. Honestly, there's no accounting for taste.

By the end of the year, both of us had reached our respective breaking points with our spouses. The marriages were over, but the geography remained – she was halfway across the country. Luckily, we already had a shared activity that put us in the same space a few times a week, which is more than most people get at the start of something new.

Our first date was appropriately nerdy. One evening after the main raid concluded, she whispered me, "hey, someone's going to show a group of us how to clip into Hyjal, wanna come?"

As if she needed to ask.

We wall-jumped up the mountains in Frostwhisper Gorge in Winterspring, and into Hyjal proper. We rode up the mountain, swam in the pond beneath the remnants of Nordrassil. One by one our friends logged off for the night, until it was just the two of us. We rode back down a ways from the summit, down to a ruined Kaldorei tower jutting out over a cliff above Felwood, and we sat there and talked until the moon set in the west across the Veiled Sea.

It was, as first dates go, perfect.

And along the way I learned something I should never have forgotten: when you want something, you go for it. Even when it's hard. Even when it's improbable. Even when the reasonable calculation suggests you shouldn't. To do otherwise isn't madness – it's just the first step down a long road that ends in a dead-end.

Within three months she had bought a house here, and we began a life together. In another year, we would be married. And today, almost five and a half years later, I can say: she's the best thing that ever happened to me.

My company email stopped working on a Tuesday.

No phone call. No meeting invite. No warning. Just a password prompt that wouldn't clear, and then a message from the contracting agency – not my boss, the *agency* – letting me know I was done. Three months into a contract. Three days before a flight to Anaheim.

We had planned the trip months in advance, because that's what you do. You book the flights before you know the lineup. You secure the hotel before the block sells out. You make plans to meet people you only see once a year, people who speak the same language and make you feel, for a long weekend, like you are among your people.

I moved through the first day like a man wearing himself as a costume. Friends tried to cheer me up. I let them try. The convention floor buzzed with the kind of energy that used to lift me automatically, and I felt it the way you feel sunlight through a window – present, visible, but not quite reaching.

We ended up in the *Diablo* hall because that's where we found seats. I was on the floor, craning up at a projected screen the size of a small building, waiting for the Warcraft keynote to begin.

I should be honest about what I had actually lost. It certainly wasn't my dream job. It was remote in the worst sense of the word. My boss worked in a different city. I had met her once in person. The work was

getting done, but at a cost that exceeded what the work was worth. I didn't enjoy it, and when it ended, my boss couldn't even be bothered to tell me herself.

Yet still, I grieved the loss of it. Not because I wanted it back, but because losing it confirmed something I had been trying not face directly: that I had drifted so far from work that mattered to me that I simply couldn't tell the difference anymore; that I was not even sure I knew who I was anymore, at least professionally; that now I would have no choice but to do so.

And then Chris Metzen walked out onto the stage.

I have been playing Blizzard games since *Warcraft: Orcs and Humans*. That is not a casual statement. It is a measure of time – of thirty-one formative years, of a specific kind of imagination, of the particular joy that comes from discovering that someone else built a world you wished you could have invented yourself.

Metzen was part of that world from the beginning. His art was in the manuals. His voice was in the games. I remember the first time I met him at BlizzCon 2007, briefly, at the signing table. He was exactly what you would want him to be: warm, present, genuinely pleased to be there. They say never to meet your heroes. They haven't met Metzen.

When he left Blizzard in 2016, citing burnout and a need to be home with his family, I was sad because I knew it was probably the best decision for him, even though it would come at a cost to both him and the rest of us, for him to be separated from something he loved.

The convention hall had not expected him back. I could see it on the faces around me – a suspended moment of collective disbelief before the room erupted. I was sitting on the floor, exhausted and embarrassed and grief-adjacent in a way I couldn't quite name.

And in an instant, all that was gone. Wiped away, albeit temporarily. It simply didn't matter anymore.

I believe that when Metzen talks about *Warcraft*, he is not performing enthusiasm. He is not simply hitting his marks, going through the motions, or just representing a brand. He is a man entirely in his element – someone who has found the exact thing he was made to do, is doing it, knows he's absolutely killing it, and cannot help but share that with everyone around him.

I recognized that feeling. I *used* to have that. I knew exactly what it felt like, and I knew that I had lost it – not all at once, not dramatically, but slowly, the way a room gets darker as the sun moves across the sky. You don't notice until you're sitting in the dim half-light, wondering when it happened.

I had let my career drift. I had taken the contracts that were available instead of the ones that mattered. I had optimized for stability and gotten neither stability nor meaning. I had become, slowly and without deciding to, someone who worked *at* things instead of *for* them.

And here was a man who had stepped away from all of it, burned down to nothing, come back – and was *lit*.

I wanted that. Not his job. Not his franchise. Not his résumé. I wanted to stand in a room and talk about my work the way he was talking about his. I wanted the feeling back.

BlizzCon 2023 was, by any conventional measure, a diminished thing. The third-party vendors were gone. The micro-experiences, the quests, the density of discovery that had characterized every prior convention I'd attended – gone. The swag bag had been replaced with an empty backpack, which I can only

describe as a metaphor so on-the-nose that Blizzard must not have noticed they were making it. The convention had stayed the same size but hollowed out, uncertain of what it was supposed to be anymore.

I understood the feeling completely.

And yet, in the *Diablo* hall, on the floor, craning my neck up at the screen, something that had been banked low for a very long time caught. Not a blaze. Not the electric shock of 2007, that first contact with something you didn't know you needed. Something quieter and more durable than that. An ember, finding oxygen.

And in that moment, a question popped into my mind, fully formed.

"What would Metzen do?"

It sounds like a joke. It isn't one. It is a genuine heuristic – a north star I return to when I'm being asked to choose between the work that pays and the work that matters. When the reasonable calculation suggests the safer path. When the odds don't favor the leap.

He walked away from the biggest job in gaming because he needed to. He came back when he was ready. He showed up to a convention full of people who love what he loves and he was *present*, fully, without apology.

Everywhere I look, I am reminded of a life lived in and around these games: in the worlds Blizzard built, in the communities that formed around them, in the friendships and the rivalries and the late nights and the convention floors and the guild raids and the first dates conducted on a mountainside that exists only in a server somewhere.

It is a life I nearly let go of. It is a life I chose, in the end, to carry forward.

This is not only a remembrance. Memory is where we've been; it is not, by itself, a destination. I am no longer the twenty-two year old who flew red-eye to Anaheim with a thin wallet and a twinkle in his eye, though I certainly recognize him. I am not the thirty-six year old packing boxes in a quiet house, though I remember exactly what that felt like.

I am something that those two people, and all the versions between them, were in the process of becoming. It's time to find out what that is.

BlizzCon 2026 is coming in September. You already know I'll be there.

The Silver Hand

"Aeon."

The name is spoken calmly, which belies the urgency behind it. We are deep into the final phase of the Lei Shen encounter, and the boss is casting Thunderstruck on the raid. Most of our raiders are badly positioned, and this is going to be a rough one. My raid leader has said my name because he has made a calculation, in the half-second available to him, that there is one person on this roster who can close the gap between what is about to happen and what should happen instead.

He doesn't know that I've already pressed the necessary button. I cast Devotion Aura a half-second before he spoke.

The neural impulse that carried my name from his brain to his mouth started at the same moment I was already in motion – he was still processing that something needed to happen while I was processing that it was *already done*. The raid will take less damage. The fight will continue. Nobody needs to know why.

This is what it means to be a paladin.

I am a paladin. That's the sentence I want to start with, because it is both very simple to say and complex to unpack.

I have played every class in this game. In the *Classic* timeline alone, I have maintained multiple characters at the raiding level across different raid teams, and at various points across that timeline have held what you might call a "main" on a rogue, a hunter, and a paladin simultaneously, distributing them across different groups based on what each team needed. I know the rogue toolkit the way a craftsman knows a blade he's carried for years. I know the hunter's rhythms, the movement, the patience and the pivot. I understand how these classes work not just as kits but as *identities* – as ways of moving through a fight and through a world.

Playing many classes is the best way to understand any one class. Often, we assume that the classes are interchangeable from the outside in the same way that roles on a team look interchangeable – somebody deals damage, somebody absorbs it, somebody puts it back. But the *how* of it differs in ways that compound over thousands of decisions. A ret paladin and an arms warrior both wear plate armor, both swing two-handed weapons, both want many of the same stats. But *they are not the same*. The warrior comes alive in the final moments of a fight, cleaving through weakened enemies during execute phase. The paladin front-loads, stacks cooldowns, erupts in a thirty-second window of concentrated devastation and then sustains through the back half with steady, deliberate output. Different shapes of the same role. Different temperaments. And where most raid teams struggle is assuming that they are fungible assets, to be swapped at a whim.

You learn this by doing it. And doing it across classes teaches you, eventually, that the one you keep returning to is not a preference. It is a reflection of your identity, expressed through the medium of a video game class.

I am a paladin.

The first paladin I ever played was in *Final Fantasy IV.* The main character is Cecil, who begins the game as a dark knight, and is forced to travel a long road of penance and arrives, on the other side of it, in white armor. His class changes to paladin – a word with which I was unfamiliar at the age of 9 – and I understood it generally as a warrior who heals. The framing made sense to me before I could have told you why. Then came Agrias in *Final Fantasy Tactics*, a holy knight in title but paladin in function – a soldier who fights to protect and defend the princess. Soon after came the paladin units in *Warcraft II*, cavalry with access to holy magic, built as much for defending allies as for breaking enemies.

The shape was consistent before the name was. Fighter and healer and defender, collapsed into one. *Versatile* – not in the sense of doing many things adequately, but in the sense of doing what the situation requires, which is a different and more useful thing.

World of Warcraft deepened all of it. *Diablo II* had introduced me to auras – the idea that a paladin's mere presence could alter the battlefield, that proximity to your allies was itself a form of service. You didn't have to activate it. You didn't have to aim it. You just had to be there, committed, and everyone around you was better for it. That struck something. The pattern continued with *Warcraft III* – the Paladin hero has an ability called Devotion Aura, which confers additional armor to every friendly unit in range. That aura has since been deprecated in *WoW* as a passive mechanic, replaced by an active ability – that through sheer force of will, you can briefly protect your raid group from a portion of incoming damage.

But I'm getting ahead of the story.

My paladin history in *World of Warcraft* begins properly at the start of *The Burning Crusade*, when I leveled one on Korgath and found myself initially raiding with it through T4 and T5 content as a healer, then into T6 as a protection paladin. The tank who built around spellpower. The one who held threat on waves of enemies in Mount Hyjal through a combination of Consecrate and sheer mechanical stubbornness, holding waves of undead on the front of my shield as if it were made of superpowered fly paper.

But if I'm honest – and this piece requires honesty – the paladin was already my answer before I had finished the question. There is a version of class selection that is purely mechanical: this class does this job, the job is needed, therefore this class. I've done that version. I've played specs and classes I'd never have chosen freely because a team needed the role, and I don't regret it, because playing things you wouldn't have chosen is how you learn what choosing means. But paladin was never a strategic deployment. It was a homecoming I kept making before I understood that's what it was.

In *Wrath Classic*, I main-tanked for Starcaller – the raid team Clover and I built from scratch – on a protection paladin. I needed that for a very specific set of reasons: I needed a class with the versatility and difference-making capabilities of a paladin, and I needed to play a class that I had thoroughly internalized the heuristics so that I could turn my attention outward to the raid. In my two other teams where I played, I was a rogue and a hunter. The three classes have lived side by side in my roster across every expansion since: rogue, hunter, paladin. I think of them as a triad, three expressions of something coherent underneath. The rogue is planning and precision – the player who has already decided what's going to happen before the pull even starts, and who suffers most when circumstances deviate from the plan. The hunter is mobility and awareness – wide-field, adaptable, comfortable in the space between structure and chaos. The paladin is the one who stands between everyone else and whatever's coming.

Paladin is the tip of the spear. The other two are what makes the spear possible.

Right now, in *Mists of Pandaria Classic*, I have four paladins at the raiding level. Three of them run in different raid groups, covering different nights, different teams. The fourth exists as a laboratory – a character where I can try things, test adjustments, see what improves before propagating a change to the others. I'm not sure what to call this except dedication expressed in the only way that makes sense at this level of the game.

I am a paladin.

Some people are happy to play a class in *WoW* because it's "meta" or because the playstyle feels good, but that's not enough for me. I love paladin because it is a class defined by its origin story.

After the First War – after the Horde tore through the kingdom of Stormwind and left it in ruins – Archbishop Alonsus Faol of the Church of the Holy Light looked at what remained and saw a problem. His priests were healers, scholars, cloth-wearers. They were spiritually formidable, but physically unsuited to the kind of warfare that was coming. What the Alliance needed was something the world of Azeroth did not yet have: people who could fight and heal and defend simultaneously, who felt the calling of the Light as something that demanded action, not contemplation.

The Order of the Silver Hand was his answer.

Warriors who saw responsibility and felt obliged to answer it. Not warriors who wanted glory. Not warriors who sought revenge for what the Horde had done to Stormwind – though the wounds were real, and the grief was real, and the anger was understandable. The best of them – Uther Lightbringer, Turalyon, those who defined what the class was meant to be – fought because fighting was what the moment required in service of something larger than themselves. To defend. To protect. To pay, in sweat and sacrifice, whatever cost was necessary to make sure the people behind them were still standing when the fight was over.

This is why Arthas is a cautionary tale rather than a disqualification.

He is the Silver Hand's greatest failure and most instructive lesson. A paladin who let his grief and his fury rewrite his sense of what he was fighting *for*. Who substituted vengeance for purpose and found, at the end of that road, something that bore no resemblance to the man who started walking it. The Order doesn't stop mattering because Arthas fell. The Order means more, actually, because it shows what the standard is and what it costs to abandon it. Uther died before he would follow Arthas into Stratholme. Turalyon spent decades in the void between worlds, holding the line, and came back. Yrel took up the mantle on a different world, fighting because her people needed someone to.

The class produces a particular kind of person. Or perhaps it attracts them.

I am a paladin.

Gap Year

The flight home from BlizzCon 2023 takes about three hours.

I spent most of it staring out the expanse of sky, thinking about the difference between grief and relief, and how sometimes you can't tell them apart. My contract was still gone, the job I hadn't loved was gone. Yet something had changed, a latent desire in me, awoken by my time at BlizzCon.

I took a few weeks when I got back. The period between losing a job and commencing the machinery of finding the next one, that's the hardest part. Between all the extra time on your hands and the lack of external feedback, you are finally alone with the questions you've been outrunning: "What did I actually want to be doing? What had I been doing, and why had it felt so hollow? Where do I go from here?"

I talked to people. I read books. I engaged a career coach. All things one should do.

The first diagnosis was technical. I had drifted so far into the management layers – PM, program leadership, communications, coordination – that I had lost the thread of actually *making* things. My skills had become almost entirely indirect. I could organize the people who built things. I could not build things myself, not anymore, not in any language anyone was currently speaking. So I went back to school, in the most informal sense: Python, Java, JavaScript, HTML/CSS, Git, the whole modern software development stack, approached with the intensity of someone who has something to prove and the advantage of having done all of this before (even if that was long enough ago that almost none of it was familiar).

I spent months in bootcamps and intensive study. I got genuinely comfortable. I felt, for the first time in a while, like I was working rather than just managing.

Nobody wanted to hire me as a developer.

A short-term project management contract appeared, and I took it – not because it was what I was looking for, but because survival doesn't negotiate. It was a temporary job in the oldest sense of the phrase: something to do while you figure out what you're actually going to do. It kept the lights on and gave me something to point to on a calendar, and when it ended, I was almost relieved. I had not been pretending, exactly. But I had not been fully myself.

By late 2024 I was a year into the gap, and I had upgraded my approach from *job search* to something harder to name – a search for meaning operating in parallel with a search for employment. I didn't just want *a* job. I wanted the *right* job.

The reading list from that period tells the story better than I can narrate it. Michael Lopp on engineering culture. David Marquet on ownership and authority. Daniel Coyle on what makes teams work and what secretly kills them. And then, gradually, a drift toward product management – toward the work of understanding *why* you're building something, not just whether you're building it correctly. Adam Grant teaching me to hold my assumptions more loosely. Simon Sinek three times, from different angles, all arriving at the same place: *purpose is not a luxury; it is structural.*

What I was diagnosing, across all of it, was an absence. I had spent years implementing things without ever being invited to ask what they meant. Off-the-shelf systems, applied because the project charter said so, handed off to a team that would hand them off to another. The work got done, but the work didn't matter – not to me, and I suspect not to most of the people in the chain. I had been a skilled functionary in systems that had no particular interest in whether I cared. And the truth, the inconvenient truth I'd been avoiding for

years, is that I care enormously. About the why. About the customer. About whether the thing we built was worth building in the first place.

That realization sharpened something.

If nobody was going to give me a context in which to work that way – and for the better part of a year, nobody had – then I was going to build one myself. That's where Aeoncore[21] came from, which gets its own chapter and deserves it. The short version: a solo-built, production-disciplined personal AI and services platform, designed to answer the question of what real operational rigor looks like when there's no team to hide behind and no budget to buy your way out of hard problems. It wasn't a portfolio project in the cynical sense. It was the only available arena.

I wish I could tell you the gap year was restful. It wasn't. It was, in many ways, harder than having a job – the functional equivalent of a volunteer position requiring full professional effort, with no compensation, no feedback loop, no external validation, and the persistent ambient anxiety of watching the clock run without knowing when it stops. There is a particular stress that comes from doing everything right and remaining unhired, that a bad job simply cannot replicate.

I am sharper than I was. More read, more practiced, more honest about what I want and what I'm capable of. The gap did what gaps, at their best, are supposed to do: it cleared the noise and left the signal. I know what kind of work I'm looking for. I know what kind of organization I need to be part of. I know, more precisely than I ever have, what I would bring to it.

I also know what I was watching, in that convention hall in Anaheim, when Metzen walked back out onto a stage he'd left seven years earlier.

I was watching someone who had found the thing he was supposed to be doing, return to doing it.

I've been working toward that moment ever since.

[21] See later chapter, "Case Study – Aeoncore Personal AI/Services Platform"

THEORYCRAFT

Systems

In the beginning was a Thing. And this Thing existed on its own, it had its own specific quality, and it was good, but it was ultimately just a Thing.

Then came a second Thing. These Things were not connected, they each had their own specific qualities, and they were good, but they were ultimately just two Things.

And then, these two Things became connected – a relationship was established, they became joined by a shared purpose – and a System was created. So while the Things remained Things, they also became components of an interrelated System. And the System was also a Thing, it had its own specific qualities, and it was good.

And so it goes.

Things are Things. Systems are a constellation of Things that are related or connected. And some Systems are in turn made up of other Systems that are also related or connected.

A hammer and a nail, sitting on opposite ends of a workbench, are not a system. They are objects. They become something more only when a relationship exists between them – when one acts on the other, when the output of one becomes the input of another, when they are arranged in a way that produces something neither could produce alone.

That relationship is the system. Not the parts. The parts are just parts.

This seems obvious stated plainly, and it is. But its implications are not obvious at all, because they cascade in both directions. Downward: any Thing, examined closely enough, reveals itself to be a System. Upward: any System, viewed from far enough away, is just a Thing inside a larger System.

Where you stand determines what looks like a Thing and what looks like a System. Both descriptions are always simultaneously true.

This has a practical consequence, and it is this: the level you're standing at limits what you can see.

Stand close enough and all you see is the Thing in front of you. That's not a failure of intelligence – it's often the job. The world runs on people who can narrow their gaze to a single object and drive it forward without being distracted by everything around it. That capacity is real, and it is valuable, and it is not what this essay is about.

Zoom out one level and the Thing becomes part of a pattern. You can see its relationships – what it depends on, what depends on it, where it fits in the sequence of things that surround it. You can start to ask not just *how do I move this Thing* but *why does this Thing exist, and what does the System around it actually need?*

That second question is the beginning of systems thinking.

It is not a framework invented in a business school and given a name so it could be put on a résumé. It is a description of something people do naturally when they've spent enough time watching local decisions produce systemic failures. When they've seen enough projects deliver perfectly against their own metrics while the program around them collapsed. When they've learned, by experience, that the question *"is this Thing healthy"* is always incomplete without the follow-up: *"healthy relative to what?"*

Nothing exists in isolation. Everything has context. Systems thinking is simply the discipline of refusing to let that context disappear.

Consider the professional stack.

An individual contributor focuses on a Thing. A task, a module, a deliverable. Their job is to execute it well. A project manager focuses on a collection of Things – the project is the System, and their job is to understand how those Things relate to each other, where the dependencies live, what the critical path looks like, and how to move the whole from here to done.

A program manager focuses on a collection of projects. The program is the System, and the job is no longer to run any individual project – it is to understand how the projects relate to each other. Where they compete for resources. Where the sequencing creates risk. How the health of the whole differs from – and sometimes conflicts with – the health of any one part.

A product manager holds something different again. Product isn't simply the next rung on the same ladder. It adds a third dimension to what had been a two-dimensional view. The project and program stack operates on the axis of execution – *how* and *when*. Product adds the axis of *what* and *why*, and it runs perpendicular to execution, not above it. You can optimize execution flawlessly on the wrong thing. The execution axis will never show you that. You need the third dimension to see it.

Beyond product: the company. Beyond the company: the market, the industry, the competitive landscape, the forces that have nothing to do with how well you built the Thing and everything to do with whether the world was ready for it. History is full of coherent products that arrived into incoherent conditions and disappeared.

The practical limit of this outward zoom is not philosophical – it is relational. You keep zooming until the System you're looking at is no longer meaningfully affected by the decision you're making, and no longer meaningfully affects it in return. The relationship has to run in both directions. When it doesn't, you've found the edge of the relevant System, and you stop there.

Not because the larger System doesn't exist. Because it isn't in play.

The same framework that describes how organizations work describes how games are built.

A game is a collection of Systems. Combat, economy, progression, narrative, social structure – each of these is its own domain of craft, with its own logic, its own practitioners, its own standards of quality. But none of them exists in isolation. A beautifully designed class that is too powerful relative to its peers is not good design. A compelling quest that cannot be completed by the character the game encouraged you to build is not good design either. The pieces have to be individually good, but more importantly, they have to support each other.

When the systems thinking in a game is strong, it disappears. The player doesn't notice the encounter was tuned against the class's damage curve, or that the quest rewards were calibrated to the economy, or that the narrative pacing was designed around the rhythm of the gear progression. They notice none of it – because

it all works. What they experience is a game that *feels right*. That flows. That makes its decisions feel meaningful and its world feel coherent.

When the systems thinking breaks down, it becomes the *only thing you notice*.

A single mis-calibrated system poisons the experience of every system it touches. The player cannot enjoy the class fantasy because the encounter design makes it irrelevant. They cannot appreciate the quest writing because the combat makes it miserable. The craftsmanship in the individual pieces is real. It does not matter, because the systems are working against each other instead of with each other.

World of Warcraft, across twenty years and more than a dozen major releases, has provided as thorough a laboratory for this dynamic as game design has ever produced. Its history is a record of systems thinking succeeding and failing at scale, in public, with millions of players providing immediate and unambiguous feedback.

Two examples from that history are worth examining at some length.

Cataclysm introduced reforging.

The mechanic was straightforward: visit an NPC, select a piece of gear, choose one of its secondary stats, and redistribute forty percent of it into a different stat not already present on the item. An item with one hundred points of hit rating and one hundred points of critical strike could become an item with sixty hit, one hundred crit, and forty haste. The worst-statted item in the loot table could be nudged toward something useful. The gap between a good drop and a bad drop narrowed.

This was the intent. An upgrade should feel like an upgrade. Even if the new item wasn't perfectly statted for your character, as long as it had more total stats and you could reforge toward your preferred configuration, the math would work out in your favor. Simplify the decision. Reduce the frustration of itemization. Give players agency over their own gear.

As a Thing, reforging was an undisputed win. Zoom out one level, and the flaw appears.

Secondary stats in *World of Warcraft* are not all equivalent in structure. Haste, critical strike, and mastery operate on curves – more is generally better, the work as multipliers in concert with one another, and the returns are relatively smooth. Hit and expertise ratings are different. Characters have a base chance to miss their attacks. Hit and expertise reduce that chance. At a specific numerical threshold – the cap – that chance reaches zero, and beyond that point, the stats stop doing anything at all. One point below the cap and they're valuable. One point above it and they are worth exactly zero.

Into this environment, reforging arrived.

The system designed to give players flexibility instead gave them an obligation. Every serious player needed to reach the hit and expertise caps before any other stat consideration was relevant. Every piece of gear had to be evaluated not for its total stats, but for how much dead-stat overhead it carried and how much reforging would be required to work around it. A new item with desirable secondaries but excess hit rating didn't feel like an upgrade. It felt like a puzzle.

Most players cannot calculate optimal reforges across sixteen gear slots in their heads, so the player base built tools to do it for them. Ask Mr. Robot arrived first. ReforgeLite followed. External simulation tools like WowSims entered the workflow for the most committed players, who would calculate their optimal

reforge configuration before even logging in. The game had externalized a core function of its own itemization system into third-party software.

The overhead was real. The fun was not.

Here is the diagnosis that only becomes visible in hindsight: reforging was not the disease. It was the immune response. Hit and expertise were already poisoning the itemization system – an item that pushed you past the cap on either stat was dead on arrival regardless of what else it offered. If you were already capped on hit and a new item dropped with hit and mastery in a slot where you wore haste and crit, it didn't matter how desirable mastery was. The hit made the item worthless. Reforging existed to minimize the damage that hit and expertise were already doing.

Blizzard removed reforging after *Mists of Pandaria*. They also – in the same patch – removed hit and expertise entirely. The decision has generally been read as a simplification of the gear system. I think it was an overcorrection, that reforging had value independent of the hit/expertise optimization. They cured the disease and the immune response went with it.

The *Mists of Pandaria* legendary questline is a more complex case, and a more instructive one, because the design intent was genuinely sound.

Legendary items had existed in *World of Warcraft* since the beginning. The model was consistent: extraordinary effort, significant exclusion. Sulfuras required crafting a powerful weapon and combining it with a rare drop. Thunderfury required two unique low-chance drops from the same boss, plus expensive materials to complete. The Warglaives of Azzinoth were purely dependent on random chance from a single end-boss. Fangs of the Father, the legendary daggers from *Cataclysm*, was the most elaborate yet – multiple solo instances requiring genuine mechanical skill, weeks of raid drops, and guild coordination to complete. It was also exclusively available to rogues. Eight of the nine available classes could not participate at all.

The pattern held across the game's history: legendary items were tremendous individual achievements, marks of trust from a guild, demonstrations of sustained commitment and skill. They were also, consistently, available only to some.

The *Mists of Pandaria* design team looked at this history and made a reasonable diagnosis, that the exclusion was a problem. And the solution they built was ambitious: A questline spanning the entire expansion, available to every class, progressing through multiple phases aligned with the content release cadence. In phase one, you collected drops from raid bosses and received a powerful gem. In phase two, player-versus-player participation earned a weapon enhancement. In phases three and four, further collection rewarded progressively more powerful items. And by the end, the final reward – a cloak, significantly more powerful than anything else available in that slot – was waiting at the end for anyone willing to put in the time.

Sustained engagement. Universal access. Expansion-length investment. The diagnosis was right, but the prescription failed to model what success would actually look like.

The player-versus-player stage is the sharpest example. PVE and PVP are not simply different activities in *WoW*. They are different *player populations*, with different motivations, different skill sets, and a long history of mutual frustration. Requiring PVE players to complete PVP content to advance a questline didn't

create meaningful cross-pollination. It created mandatory participation in an activity many players actively disliked, with no path around it and no accommodation for the skill gap. And it created a second failure simultaneously: PVP players, who had built their own ecosystems of skill and competition, found their battlegrounds suddenly flooded with inexperienced players who didn't want to be there. The stage created resentment in both directions at once.

The catch-up problem compounded the pressure. A player who missed a phase didn't simply fall behind – they had to complete the previous phase's requirements before they could begin the current one. The questline penalized absence in a commitment designed to span an entire expansion. What was intended as an inclusive alternative to guild-gated legendaries became its own form of compulsion: you had to be there, consistently, or the cost of returning grew with every week away.

The per-character design added another layer. In a game where many players maintain multiple characters, the questline created a standing obligation for each of them. Even casual attachment to an alternate character meant feeling compelled to keep its questline current – not because you were certain you'd play it, but because the cost of falling behind made the decision feel irreversible.

And then, the detail that is almost funny in retrospect: the final raid tier of *Mists of Pandaria* contains boss drops in the cloak slot. The questline rewards a legendary cloak. The raid drops non-legendary cloaks. For any raid group keeping pace with the questline – which the design explicitly encouraged – those boss drops were obsolete on arrival. Players would see a cloak appear in the loot window and immediately become disgusted at the "wasted loot." The system had not accounted for what the loot table would feel like once its own questline succeeded.

The failure here is not in the ambition. It is in the modeling. The designers asked *"how do we make everyone feel included"* and built a System to answer that question. They did not ask *"what does this system look like when it works, and what does the game feel like then?"* The questline succeeded on its own terms. The experience it produced was not the experience that had been intended.

Building a System and reading a System are the same cognitive act, run in opposite directions.

The designer asks: how do these Things need to relate to produce this experience? The player – the player paying the right kind of attention – asks: what relationships must have been designed here to produce what I'm experiencing? Same map. Different starting point. Same destination: an understanding of the whole that cannot be reached by examining any single part.

Thirty years of playing games at high levels of engagement and competitive investment is a long time to practice asking that second question. Long enough that it stops feeling like analysis and starts feeling like perception. The systems become visible the way grammar becomes visible to a writer – not as rules being consciously applied, but as structure that is simply, naturally, there.

What you learn to see in a game, you learn to see everywhere.

The project that is late for reasons the project manager cannot explain because they cannot see the program around it. The product that is technically excellent and commercially irrelevant because no one modeled the market it was entering. The organization optimizing every individual metric while the system

itself fails. These are not different problems. They are the same problem, wearing different clothes, at different levels of the same stack.

The only sane response – in a raid, in a program, in a product organization – is to keep asking what system this Thing belongs to, and what that system actually needs.

The answer is almost never what the local view suggests.

Agility

In 2001, while my friends and I were trying to work out the finer points of the Zergling rush against one another, another group of people were thinking about a better way of building software. They met at a ski lodge in Utah, and over a long weekend, wrote the Agile Manifesto[22].

Specifically, it was about the growing chasm between how software was being built and what the people paying for it actually needed. The prevailing model – now called Waterfall, though nobody called it that at the time – worked like this: define the requirements, write them all down, create a detailed plan for how to deliver the project, then follow that plan exactly. By the time the thing was finished, months or years had elapsed, the world had moved, and what had been built was often an answer to a question nobody was asking anymore.

The seventeen people in the ski lodge were practitioners. They had all, in various ways, been trying to build software differently – iteratively, collaboratively, with feedback loops short enough to actually learn from. They each had different names for what they were doing, different emphases, different flavors. But the underlying conviction was the same: that a plan is a hypothesis, not a contract, and that clinging to it past the point of relevance is not discipline. It's denial.

They wrote sixty-eight words (it's shockingly short). Called it the Agile Manifesto. Released their work on a world that seemed to need to hear it.

The revolution began, but within a decade, Agile had become what it was invented to replace. It wasn't that the words themselves changed, but what they signified in practice. Certification programs appeared. Then certification bodies. Then consultants who certified other consultants. Books arrived by the thousands, each one explaining, expanding, and elaborating on a manifesto that had been written precisely because too many people were explaining and expanding and elaborating instead of building. Ceremonies emerged: the daily standup, the sprint retrospective, the sprint planning session, the sprint review, the backlog grooming session, the story point estimation, the definition of done. All of it sounded so reasonable, and in the right proportion, maybe it was.

And yet, all of it changed what Agile meant in practice, to the point that it became exactly the thing the Agile mindset was supposed to prevent.

The irony is tidy enough to hurt: a philosophy built to minimize busywork became a cult with its own dogma and prescriptions. The question "*are we building the right thing?*" got replaced by the question "*are we doing this correctly?*"

And that replacement – the drift from outcome to ritual – is where most Agile implementations go wrong. Not in a way that is catastrophically obvious, but in a way that people eventually look up from their desks and think "how the heck did we get here?"

Strip the barnacles from the hull, and the manifesto says four things. People matter more than the processes that coordinate them. The product matters more than the documentation that describes it. Customer outcomes matter more than contracts that specify deliverables. Responding to change matters more than adhering to a plan. The original document adds: "*while there is value in the items on the right, we value*

[22] See "Manifesto for Agile Software Development" (https://agilemanifesto.org/)

the items on the left more." That qualifier is doing significant work. This is not an argument for chaos. It is an argument for priority.

Build. When the world shifts, shift with it.

In 2017, a musician named Damian Kulash (of the band OK Go) gave a TED talk about how the group comes up with their music videos, entitled "How to find a wonderful idea."[23]

OK Go is not famous for its music – or rather, it's more appropriate to say that they are considerably more famous for their wildly creative music videos. The treadmill video. The zero-gravity video. The Rube Goldberg video, which required 89 takes and a warehouse full of machines built from ordinary objects that had to work in precise mechanical sequence. The work is elaborate, technically ambitious, and resists easy categorization. The question everyone asks is: where do these ideas come from?

Kulash's answer is not what you would expect. He notes that when we think of looking for a creative idea, what we're really looking for is a sense of wonder, and that an inherent quality of that wonder is a sense of surprise. To get there, he does not describe a creative process. He describes the absence of one – or rather, he describes why the conventional creative process is structurally incapable of producing the kind of ideas he is looking for.

The conventional process, as he tells it, goes like this. You think of an idea. You make a plan to execute it. You check each step off as you complete it. (This should sound very familiar to waterfall practitioners.)

The logic is sound, but the math is punishing: in a sufficiently complex project, even with 99% reliability at each individual task, the accumulated probability of something going wrong approaches certainty[24]. Which means, as Kulash observes, that *"if your project is pretty complex... you're basically constrained to just reshuffling ideas that have already demonstrably proven that they're 100 percent reliable."* You cannot plan your way to a surprising idea, because planning selects against surprise. The only ideas safe enough to commit to in advance are the ones that have already been done.

His solution is a sandbox.

"What we do is we try to identify some place where there might just be a ton of those untried ideas. We try to find a sandbox and then we gamble a whole bunch of our resources on getting in that sandbox and playing."

Play, here, is not frivolity. It is a specific discipline – one that deliberately suspends judgment long enough to let unexpected combinations surface. The sandbox is full of seeds: partial ideas, interesting materials, unexpected juxtapositions, techniques that haven't found their purpose yet. You don't plan which seed will become something. You can't know. You put the right people in a space full of the right raw material and you let the process reveal what's possible.

---

[23] It's on YouTube, "How to find a wonderful idea | OK Go" (https://www.youtube.com/watch?v=WyOSqjIABe0). Absolutely worth watching.

[24] The math works out that if an individual task has a 99% reliability factor, with 100 total tasks you only have a 34% chance that all of them will have been successful. In other words, you reach the point where you're twice as likely to fail than succeed.

And when it works – when the right combination of ideas finds its shape – the experience is distinctive. *"When it does click, it doesn't feel like you thought up that puzzle piece, it feels like you found it – like it was a set of relationships that you unlocked."*

Kulash never uses the word Agile (after all, he's a musician, not a PM). He is describing a band's creative process, not a software methodology. Even so, the resemblance is not incidental.

The delivery half of game development is easy to recognize as Agile territory. Sprints. Milestones. Iterative builds. Ship a slice of the thing, see what breaks, adjust, ship again. The vocabulary maps directly. Anyone who has run a release cycle in an enterprise software environment would recognize the cadence.

The creative half is where it gets interesting – and where most organizations either don't look, or don't think to apply the same logic.

Worldbuilding is a sandbox. Narrative design is a sandbox. Class design, encounter design, economic tuning – all of these are fundamentally iterative disciplines. You do not design a compelling villain by specifying requirements and executing against them. You put ideas in conversation with each other until something surprising emerges that is also, somehow, exactly right. I wasn't in the room, but I cannot believe that Arthas was the product of a planning session. He was a set of relationships that unlocked. The long arc – from prince to paladin to death knight to Lich King – has the feeling of inevitability in retrospect, the way found things always do. But it does not feel engineered. It feels discovered.

The organizations that do this well have internalized one thing: the process that reveals good ideas and the process that ships reliable software are not two different processes running on separate tracks. They are the same cognitive posture, expressed at different scales, in different domains. Build small. Observe what emerges. Let reality reshape the next iteration. Trust that the sandbox will show you which ideas are not only surprising, but – in Kulash's phrase – *surprisingly reliable.* The goal is not to predict the right answer. It is to create the conditions in which the right answer can appear.

When the delivery discipline and the creative discipline are aligned – when both halves of a game studio are genuinely iterating, genuinely responsive to feedback, genuinely willing to let the plan change when the work reveals something better – the result is a kind of organizational fluency that is difficult to describe and immediately recognizable when you encounter it. The game feels coherent in a way that designed things usually don't. The systems support each other. The world has internal logic that wasn't fully documented because it wasn't fully known until it was built.

When they fall out of alignment, it's the only thing you notice.

A creative team doing genuine sandbox work alongside a delivery pipeline locked into fixed-scope commitments is not Agile. It is two organizations with incompatible epistemologies sharing a building. The creatives will change their minds, because that is what the process requires. The delivery team, if it has drifted back toward ceremony, will experience those changes as failures of discipline rather than signals to be incorporated. The argument that follows is not a personality conflict, it is a philosophical one.

The Agile manifesto wasn't simply about how to make better software. It was a revolutionary idea with surprisingly universal application: that good work actually happens when people quit worrying about being procedurally correct and focus on what really matters. Build. Embrace change.

That is as true in a warehouse full of Rube Goldberg machines as it is in a sprint review. As true in a writers' room as in a backlog refinement session. The domain changes. The posture doesn't.

The ceremony was never the point.

Alchemy

"The trouble with market research is that people don't think what they feel, they don't say what they think, and they don't do what they say."

I read that sentence and thought: "damn, that's describing gamers."

I was reading a book written by Rory Sutherland, titled *Alchemy*, and the provocation is in the subtitle: *The Surprising Power of Ideas That Don't Make Sense.* Sutherland is a fascinating figure: he's an advertising wizard, a master of behavioral psychology in both theory and practice, and he has that incisive wit we typically attribute to the British.[25]

When I say that I interpreted that quote as being about gamers, I mean that, not as an insult, but as a surprisingly honest assessment. Gamers are, as a population, among the most vocal, opinionated, and self-aware consumer groups in the world. They produce prodigious quantities of feedback. Forums, subreddits, Discord servers, YouTube essays, Twitch streams running live commentary on design decisions. They know what they want. They will tell you at great length. They are often wrong about themselves in ways that would be funny if the consequences weren't so real.

Players who demand faster content delivery and will then burn out when they get it. Players who said they hated grinding will then turn around and play for eight hours straight when the grind has a good feedback loop. Players who complain about the amount of trash mobs in a raid dungeon will then complain at how empty a trash-less raid feels[26].

Sutherland's point, and Ogilvy's before him, is that the map is not the territory – and the territory, in this case, is human psychology, which does not run on logic.

Sutherland calls this *psycho-logic*. His argument is not that logic is wrong, it's that logic is incomplete. That logic describes how things should work given a rational actor making decisions in their own best interest, but human beings are not rational actors at all; we are creatures of perception, context, and feeling, and the decisions we make are downstream of those things, not borne from the objective conditions around us.

The practical implication is uncomfortable: if you optimize purely for the rational solution to the problem your users have described, you will frequently produce the wrong answer[27] – because the rational solution solves the rational problem, which is usually not the real one.

Blizzard has known this for a long time. In 2010, Rob Pardo – then EVP of game design – gave a talk at GDC in which he described a principle the studio called "Make It A Bonus." The example he used was the rest experience system in *World of Warcraft*.[28]

When *WoW* launched, players who spent extended time logged in would incur a penalty: their experience gain dropped from 100% to 50% after a few hours. The intent was clear enough: encourage

---

[25] I've had videos and talks from him pop up on my YouTube feed before, but it's only fairly recently that I sat down to read his book. It's caused me to re-think almost every facet of my career and experience, and I think everyone should read it.

[26] As they did in Trial of the Crusader from *Wrath of the Lich King*.

[27] Or, if not strictly *wrong*, it will be potentially inefficient or incomplete.

[28] A summary of Pardo's talk can be found here: (https://www.gamedeveloper.com/game-platforms/gdc-blizzard-s-core-game-design-concepts).

breaks, prevent marathon sessions, manage the rate at which players progressed through content that had taken years to build. Beta players hated it. The penalty felt like a punishment, like the game was scolding them for playing it.

The fix was simple. Instead of dropping players from 100% to 50%, the system was reframed: players who had been logged out for a period would accumulate rested experience, starting their next session at 200%. They would play at that elevated rate until the bonus was exhausted, then settle back to 100%.

In practice, nothing actually changed: you had the same relationship between time played and rate of gain. Yet it felt completely different.

The penalty became a reward. The feeling of being punished for engagement became the feeling of being rewarded for returning. Players didn't complain about the drop-off from 200% to 100% the way they'd complained about the drop-off from 100% to 50% – because 100% felt like a baseline, not a punishment. What Pardo's team had understood, and what Sutherland would later give a name, is that the nature of our attention affects the nature of our experience.

This is psycho-logic. And Blizzard was practicing it years before *Alchemy* existed as a book.

But the most interesting things Blizzard has done with alchemical thinking aren't the elegant reframes. They're the things that should not have existed at all.

Sutherland has a rule he calls "Dare to be trivial." His observation is that the most useful things – the things that actually change how people feel about a product – are frequently the ones that make no rational sense to build. They don't contribute to the core experience. They aren't defensible in a product planning meeting. They're the features that a sufficiently rigorous rational process would never have permitted, because a sufficiently rigorous rational process would have correctly identified them as frivolous.

Consider the Murloc.

The Murloc is a small, amphibious, vaguely fishy humanoid that makes a hilarious (or terrifying, depending on your situation) gurgling noise[29] and attacks in swarms near bodies of water in *World of Warcraft*. They are, by any functional measure, unremarkable: a minor enemy with a distinctive sound file. They're unimportant in the lore, contribute little to the story, and by every rational metric for evaluating the value of a game asset, they should be a footnote – one of hundreds of creature types, long forgotten.

Instead, they've become an iconic part of the game universe.

Plushies. Costumes. A running joke embedded in BlizzCon mythology. An in-game pet, then multiple pets, then a Murloc version of nearly every significant character Blizzard has ever made. The Murloc is ridiculous, and the ridiculousness is precisely the point – it is the thing that, in Sutherland's phrase, logic would never have built, and therefore only alchemy could have produced. Blizzard kept making Murloc things because players kept responding to Murloc things, and the players responded because the Murloc represented something that careful, rational, defensible game design cannot manufacture: genuine absurdist delight.

[29] "MRGRLGRLGLRGRLGRL!" – every Murloc ever

Pet Battles. The Diablo cow level. The Brawler's Guild. The various holidays and seasonal events that annually populate Azeroth with pumpkin-headed horsemen and winter festival decorations. A product planning document would have cut all of them. Players would have listed none of them when asked what they wanted. And yet these are the things people bring up when they try to explain why *World of Warcraft* felt alive in a way that other games didn't.

Dare to be trivial. Because the trivial things are often doing the most important work.

The boldest application of alchemical thinking in Blizzard's history is also the one that the market never fully rewarded – which does not mean it was wrong.

Sutherland has a number of rules for alchemical thinking, and one of the primary ones is this: "the opposite of a good idea can be another good idea."[30] By the time Blizzard was thinking about entering the MOBA genre in the mid-2010s, there were already two competitors on the scene – *League of Legends* and *DOTA 2* – and there was a fairly rigid formula to how those games worked.

Heroes of the Storm looked at every one of them and asked: what if we did the *opposite*?[31]

Multiple maps, each with different objectives – not simple variations on a theme but genuinely different strategic problems requiring different team compositions and approaches. No gold, no item shop. A talent system instead, where players chose from branching upgrades at fixed intervals, shaping their hero's identity across the course of a match without ever visiting a vendor. And shared team experience – not individual levels, but a single team level, so every player on the team was always the same strength, so there was no feeding, no feeding penalty, no single player's poor performance dragging the team's power down.

Each of these decisions looks wrong from inside the MOBA genre's rational framework. Gold and items give players agency. Individual levels create meaningful skill expression. The single map creates a known, fair, legible competitive environment. *Heroes* gave up all of it.

What it got in return was a fundamentally different emotional experience. Without feeding, every player on the team remained relevant regardless of individual performance in any given fight. Without an item shop, the cognitive overhead of build optimization – which in *League* and *DOTA* is itself a deep skill layer that some players love and others find alienating – disappeared, and what replaced it was moment-to-moment objective decision-making. Without a single map, the game never settled into the pattern-matching that characterizes high-level play in map-static games. Every game was a slightly different game.

The corollary to the 'opposite rule,' which he doesn't always say out loud but which the examples make clear, is that you will never find the opposite of a good idea by optimizing within the existing framework.[32] You can only find it by being willing to interrogate the framework itself – to ask not "how do we do this better?" but "what are we actually trying to achieve, and does the conventional approach actually achieve it?"

---

[30] Note: the distinction is not that it will generate a *better* idea, but that both can be equally good and worth seeking.

[31] For further examination of what I'm talking about, see earlier essay "Elegy of the Storm."

[32] Sutherland attributes this to personal risk, that few people are willing to stick their necks out on an idea that is the opposite of what logic suggests for fear of losing their job or position. My assertion is that bravery is an undervalued characteristic in today' business leaders.

Heroes asked those questions. Its market performance may not have been what Blizzard hoped for. As a player and spectator, though, I found it to be a superior experience to the competition.

The honest thing to say about alchemical thinking is that it doesn't come with a guarantee.

Sutherland acknowledges this with what is possibly his most useful rule, stated with characteristic cheerfulness: "A good guess which stands up to empirical observation is still science. So is a lucky accident." He is making a point about epistemology – that we demand more certainty from unconventional ideas than from conventional ones, and that this asymmetry is not rational. The rest experience system didn't emerge from a theory of psycho-logic. It emerged from listening to frustrated beta players, finding a reframe, and discovering that it worked. The discovery came first; the explanation came later.

This is how a lot of good design actually operates, and the refusal to acknowledge it is one of the ways rational frameworks produce worse outcomes than they should. We require that unconventional ideas justify themselves in advance, with mechanisms we can explain and benefits we can project. We require nothing of the sort from conventional ideas – the conventional approach is innocent until proven guilty, and the alchemical approach is guilty until proven innocent.

The practical correction is not to abandon rigor, but to apply it evenhandedly. To test the counterintuitive thing because nobody else will. To consider the opposite of the good idea because the framework, by construction, will not consider it for you. To recognize that the gap between what players say they want and what they actually respond to is not a flaw in the research. It is the most important data point in the room.

Systems thinking tells you that a Thing is only legible relative to the System it belongs to. Agile tells you that the best way to advance that system is through iteration and adaptation – that you cannot plan your way to a good answer, you can only build your way there.

Alchemy tells you that the humans inside the system are not running the decision-making process you think they are. They don't think what they feel. They don't say what they think. They don't do what they say.

All three of these are true simultaneously. None of them is sufficient on its own. Systems thinking without iteration becomes architecture without function. Agile without systems thinking becomes velocity without direction. And both of them, without an alchemical sensibility, become the most elegant possible answer to the wrong question.

The professional I want to be – and the professional I believe Blizzard needs more of at the intersection of delivery and design – is one who holds all three. Who zooms out far enough to see the system, who builds iteratively enough to let the system reveal what it actually needs, and who remains suspicious enough of rational models to ask, at every significant decision point: are we solving the problem people have, or the problem they've described?

They are not always the same problem.

Usually, they aren't.

The Infinite Game

For a period of about ten years, a particular phrase circulated through gaming press cycles with the reliability of a seasonal allergy. It appeared in preview coverage and forum threads and the comment sections of every major gaming outlet. It followed each new massively multiplayer online game to market with the confidence of a campaign slogan.

WoW killer.

The term carried a theory of competition embedded in it – the idea that the correct objective for any new entry into the MMORPG market was to defeat *World of Warcraft*. To take its players, collapse its subscriber numbers, and claim the throne. *Lord of the Rings Online* launched with the Tolkien legendarium as its entire value proposition and a player base hungry to walk the Shire. *Age of Conan* arrived with a darker, more visceral aesthetic and the full weight of Robert E. Howard's source material behind it. *Star Wars: The Old Republic* came backed by BioWare's storytelling reputation and one of the most beloved franchises in the history of popular culture. *Wildstar* – developed by a team that included former Blizzard employees who knew exactly how the machine worked – arrived swinging hard at *WoW's* casual-to-hardcore spectrum and missing both.

None of them killed anything. Most of them aren't even around anymore.

And then, against the odds, and without declaring its intentions[33], *Final Fantasy XIV* found a permanent home in the market.

Simon Sinek didn't invent the distinction between finite and infinite games[34] – that credit belongs to the philosopher James Carse, whose 1986 book put the framework on paper – but Sinek popularized it, sharpened it into a business lens, and made it useful in contexts Carse hadn't fully explored. The core concepts are simple enough to state in a paragraph, and worth stating clearly before going further.

A finite game has fixed rules, a known set of players, a defined timeframe, and a clear winning condition. You play to win, and when someone wins, the game is over. Chess. Football. A sales quarter.

An infinite game has none of those constraints. The players are known and unknown – competitors you can see and ones you can't anticipate. The rules shift. Players can enter or exit at any time. There is no winning condition, no terminal state, no final whistle. You play to keep playing. And the only way to *lose* an infinite game is to exhaust your will, your resources, or both, and drop out.

The critical implication – the one that explains a lot of corporate failure – is what happens when a player built for an infinite game starts treating it as finite instead. When they stop asking "how do we continue to improve and compete" and start asking "how do we *win*?" The target changes. The strategy narrows. And the narrowed strategy, optimized for a finishing line that doesn't exist, slowly undermines the capacity to sustain the actual game being played.

---

[33] Not to mention a painful death and rebirth due to its initially failed reception among the community.

[34] I love the concept in this context because the infinite game refers both to the business' operating model, but also the way live service games are constructed. There are finite games within *World of Warcraft*, but ultimately the game can continue as long as Blizzard has the will to keep building to it and the community wants to keep playing it.

Every game on that list of WoW killers made this mistake. Not because they were poorly made – some of them were genuinely good – but because the objective they organized around was the wrong one. *Beat World of Warcraft* is a finite objective. It implies a scoreboard, a threshold, a moment at which victory could be declared and the game concluded. The MMORPG market is not a finite game. It never was. And playing it as one – optimizing for the kill shot rather than for long-term viability – produced exactly the outcome you'd expect.

They ran out of resources, or will, or both. And they stopped playing.

World of Warcraft has been running for over twenty years. It has survived expansion cycles that ranged from generational highs to periods that even its most devoted players will acknowledge were not its finest hours. It has weathered corporate turbulence, public controversy, a global pandemic, and the sustained structural pressure of a gaming landscape that looks almost nothing like the one it launched into in 2004. It is still here.

Part of that longevity is the game itself – the accumulated weight of twenty years of lore, the embedded social infrastructure of guilds and communities and friendships that exist nowhere else, the sheer inertia of a world that millions of people have called home across multiple chapters of their lives. But part of it is something more deliberate.

World of Warcraft is a subscription product. It has always been a subscription product. The business model is structurally an infinite game: the objective is not to sell the most copies at launch, but to give players a reason to stay subscribed month after month, year after year, expansion after expansion. Every design decision that improves the long-term experience – and every decision that sacrifices it for short-term engagement metrics – shows up eventually in the subscriber numbers. The feedback loop is real, and it runs in both directions. Blizzard has occasionally felt that pressure in ways that weren't comfortable.

But the model itself is honest about what the game actually is. You are not playing to finish. You are playing to keep playing.

The players understand this, even if they've never read Sinek. The ones who have been there since Vanilla don't talk about *beating World of Warcraft*. They talk about *living* in it. The goal was never the end boss. The goal was the world.

Here is where *Final Fantasy XIV* becomes instructive, and where Sinek's framework adds something specific.

In *The Infinite Game*, Sinek argues that a worthy rival is not a threat to be neutralized – it is a necessary feature of a healthy competitive environment. The right competitor forces you to improve. It reveals your gaps. It keeps you from becoming complacent atop a position you've stopped having to defend. The worthy rival is not your enemy. It is, in a meaningful sense, part of what keeps you in the game.

FFI4 is that rival now, and it got there by refusing to be a *WoW* killer.

The early history of *FFI4* is not something I'll dwell on – it launched in a state that its own developers later acknowledged was indefensible, and the story of how it was rebuilt into *A Realm Reborn* is well documented and not the point here. The point is what it became after the rebuild, and why.

FF14 did not set out to take *WoW's* players. It set out to be the best version of a Final Fantasy MMO that it could be. The design philosophy, the narrative sensibility, the community culture that formed around it – none of it is optimized to appeal to lapsed *WoW* players who are angry about their current subscription. It is optimized to be *FF14*. And the audience it found is, to a substantial degree, an audience that might never have been comfortable in Azeroth in the first place.

The two games borrow from each other occasionally, as any healthy competitors do. *WoW* recently implemented player housing, a feature *FF14* has offered for years[35]. The cross-pollination is real, and it's evidence of exactly the dynamic Sinek describes: two players in an infinite game, each improved by the presence of the other, neither trying to deliver a finishing blow. If *FF14* shut down tomorrow, Blizzard would not issue a press release declaring victory. They would keep telling the *Worldsoul Saga*. If *WoW* went dark, *FF14* would not expand to absorb the displaced population as a strategic maneuver. It would keep being *FF14*.

That is what a healthy infinite game looks like from the outside: two players, occasionally ahead and occasionally behind, focused on their own capacity to continue rather than on eliminating the other.

The earlier competitors – *LOTRO, AoC, Wildstar*, the others – had the raw material to play the same game. They had intellectual property, talented teams, and genuine audiences who wanted what they were offering. What they didn't have, or didn't hold onto, was a clear answer to the question *why do we exist, independent of our position relative to World of Warcraft?* Without that answer, the competitive pressure collapsed their strategy inward, and they found themselves optimizing for a kill they couldn't land rather than a game they could sustain.

The worthy rival survived. The *WoW* killers didn't.

Blizzard is not immune to finite thinking. It would be dishonest to suggest otherwise, and this essay isn't a brief for hagiography.

Heroes of the Storm was, at least in my own accounting, the best MOBA ever made. The design philosophy was genuinely distinctive: no items, no last-hitting, an emphasis on teamfight and macro strategy over individual mechanical expression that made it more accessible to casual players without sacrificing depth for the dedicated ones. The community that formed around it was passionate. The esports scene was ambitious.

And in 2018, Blizzard significantly scaled back development and shut down the HGC – the Heroes Global Championship – with two weeks' notice, mid-season, devastating a professional scene that had organized their lives around it. It was a decision executed badly, and the people hurt by the abrupt shutdown deserved better handling.

But, strip away the execution and look at the underlying logic: a product that was beloved but commercially underperforming, competing in an esports ecosystem dominated by *League of Legends* and *DOTA 2* – two games with years of established infrastructure, massive viewership, and network effects that

[35] By extension, you could argue that *FF14's* "glamour" system, an appearance-modification mechanic, is clearly cribbed from *WoW*'s transmogrification system.

Heroes was unlikely to overcome regardless of design merit. The resources allocated to the HGC were finite. The question was whether they were better deployed sustaining a losing competitive position, or redirected toward *Overwatch 2* and *Diablo IV*.

From an infinite game perspective, that is actually a coherent decision. Sometimes stewardship of the whole means accepting a local loss. Sometimes the right move is to release one piece in order to strengthen the board, like sacrificing a bishop to capture a queen. The mistake was not in cutting; it was in how the cut was made. The obligation to the people who had built their professional lives around the HGC deserved more than two weeks' notice – and that's a legitimate criticism that exists independently of whether the underlying resource decision was right.

Diablo Immortal was rightly criticized. *Warcraft Rumble* came and has largely gone. These are not nothing. But they are also not structural failures – they are the variance you'd expect from an organization willing to experiment, run through a portfolio diversified enough to absorb the weight of individual misses. The overall health of the enterprise, and the goodwill of its community, remained intact because the aberrations were legible as aberrations. They stood out precisely because they weren't the norm.

The risk, and it's worth naming clearly, is cumulative. Finite-game thinking doesn't usually announce itself. It accretes. A decision to optimize this quarter's engagement metric. A feature designed to drive retention numbers rather than player experience. A roadmap shaped more by competitive anxiety than by genuine creative vision. None of these look like much in isolation. Stacked together over years, they are how you wake up atop a lonely throne, having gradually displaced the thing that earned you the throne in the first place.

Blizzard knows this. The return of Chris Metzen, the explicit recommitment to the *Worldsoul Saga*, the careful stewardship of the *WoW Classic* ecosystem – these are not the moves of an organization that has forgotten what game it's playing. They read, to me, like someone consciously pulling the organization back toward the long view. And I think that read is correct.

I am aware that I have just spent several pages applying the frameworks of a business author to a video game company, and I want to be direct about why.

The work I would do at Blizzard – whatever its specific shape – would be played in the same register. Not the infinite game of the MMORPG market, but the infinite game of a game studio committed to worlds it intends to inhabit for decades. The *Worldsoul Saga* is not a three-year roadmap; it is a generational creative commitment. The *Classic* ecosystem is not a feature; it is a parallel timeline being maintained alongside the living game. These are infinite game decisions, and they create an infinite game context for every producer, every program manager, every person responsible for keeping the delivery machine running.

In that context, optimizing for a sprint at the cost of a saga is not discipline. It is the wrong game.

This connects to something the systems thinking essay tried to establish: that finite objectives are often embedded inside infinite games, and that the skill is in knowing which level you're operating on at any given moment. An individual raid tier is a finite game – a beginning, a middle, a clearance, a conclusion. The expansion is a longer finite game, with its own arc. The franchise is something else. It doesn't end. It evolves.

And the question at that level is never *did we win this release* but *did we leave the game healthier than we found it.*

That is the question I would carry into every room.

Not *did we ship on time* – though we would, because that matters too. Not *did we hit the quarterly engagement numbers* – though we'd track those as signal, because they are. The first question, always, underneath all of it: *did we do something today that makes the long game more possible?*

You play the infinite game well by refusing to lose sight of what it actually is. The point isn't to win, the point is to keep playing.

Pathfinding

Don Norman, in *The Design of Everyday Things*, describes a forcing function as a physical or systemic constraint that prevents a certain action from occurring until a prerequisite condition is met. The car that won't start until your seatbelt is fastened, or the door that only opens in one direction – those are forcing functions. Forcing functions aren't punishments; they are design decisions, deliberately narrowing what is possible in the service of what is desirable. Norman's insight is that constraint, applied intelligently, produces better outcomes than open possibility ever could.

Which brings me to flying mounts.

Flying mounts arrived in *The Burning Crusade*, and they were immediately, obviously, undeniably wonderful. Outland had been designed with the possibility in mind, its seven zones laid out with enough geometric intentionality that players could safely inhabit three dimensions for the first time. You could hop on your gryphon, point yourself at a distant objective, and simply go – no cliffs, no roads, no wandering mobs between you and where you wanted to be. The traversal problem, which had been a part of navigating Azeroth since the dawn of *WoW*, was solved. Players were ecstatic.

Blizzard, in retrospect, was not.

The Pandora's box problem with flying mounts wasn't that they made travel easier – that was the point. The problem was what players did with that ease. Why learn the terrain when you could fly over it? Why pick your way carefully down a cliffside when you could simply flying off the edge and glide to the bottom? Players are optimizers by default, and they'll gleefully min-max almost any situation. The moment a faster path exists, it becomes the only path worth taking. And so the world – Blizzard's world, built by hundreds of artists and designers and writers – shrank to a collection of quest pins on a map. The in-between spaces, the ones that made Azeroth feel like a place rather than a backdrop, became the floor beneath your airspace.

The developers knew it. But flying mounts had become an expectation. You couldn't take them away. Players would riot in the streets.

And then, with *Warlords of Draenor*, Blizzard did it anyway.

The announcement was well-reasoned but still irritating: "We intend to disallow flying while leveling from 90 to 100, and have flying become available again in the first major patch." Players grumbled, complained, argued, and then – because the game was still *World of Warcraft* – played anyway.

And something happened in that forced return to earth. Draenor, experienced at ground level, was a different game than Outland had been. You noticed where the roads curved. You learned which cliff faces had navigable switchbacks and which ones didn't. You developed something that might be called player proprioception – a spatial awareness of your position in the world, an instinctive sense of how the terrain related to itself. It was the difference between reading a map and knowing a place so well you don't even need the map. Blizzard had, by removing the option to bypass the world, forced players back into it.

Then came Pathfinder.

Pathfinding was a meta-achievement: For each expansion zone, unlock the achievements for Explorer (discover the full map) and Loremaster (complete the major quest lines). The logic was simple: do the work, see the world, and you'll earn the right to fly above it. And once earned, it was account-wide – unlock it on one character, and every other character you owned could fly in Draenor from the moment they arrived. The

grind, such as it was, was paid once, and the dividend compounded across your entire account. This was Norman's forcing function, applied to a live game with millions of players: a prerequisite condition, deliberately designed, that you had to satisfy before the restricted action became available. The thing that felt like a removal felt, by the end, like an achievement. Players stopped complaining.

Here is where the argument turns, because the forcing function wasn't just applied to players. It was applied to the design team too.

The implicit contract of flying mounts in *TBC* and beyond had always been uncomfortable for the designers building the world: *if players are going to fly over the terrain anyway, how much does the terrain actually matter?* It's a corrosive question, and the honest answer is that it showed. Blade's Edge Mountains in *The Burning Crusade* was a zone of stunning visual ambition and punishing navigability – jagged canyon country with no clear paths and poor signposting, the kind of place that felt arbitrary rather than discovered. I avoided that zone at all costs, only returning once I could use my gryphon to make up for the poor geometry.

The guarantee that players would actually *see* the ground changed what ended up on it. *Warlords of Draenor*, designed with the expectation of ground-bound players, featured terrain built for discoverability – ramps where ramps made sense, paths that rewarded exploration rather than punishing it. Vignettes and small moments of worldbuilding filled the in-between spaces because the team knew players would pass through them, not over them.

I give *Mists of Pandaria* a lot of credit for advancing the state of making Azeroth feel real and *lived-in*. The Valley of the Four Winds remains, to this day, one of the most convincing zones in the game's history – actual farmland, actual villages, a landscape that feels inhabited and worked and aged. Blizzard went the extra mile there even knowing that most players would only experience it at ground level on their first character, with alts flying in from the start. They built it anyway. The constraint didn't fully exist yet, but you can tell that the intention was there.

By the time we reach *Legion* – the first expansion designed from the ground up within the Pathfinding framework – you can see what that intention produced at full maturity. Suramar is not just a zone. It is the argument that constraint makes. Ancient, layered, crumbling-elegant, lived-in in a way that zone design rarely achieves – every surface of it rewards attention. The team built it knowing players would be on the ground, and the ground is worth being on.

The Pandora's box problem with flying mounts was never fully solvable. You can't un-introduce a feature that players love, because you can't turn back the expectation you've created in your players. But you can reshape the conditions under which it's available – and in doing so, reshape what players and designers alike bring to the game. That's the forcing function at its most sophisticated: not punishment, not removal, but redesign. Blizzard didn't take flying away, they made you *earn* it, and along the way improved the experience and the product at the same time.

Delivery

I have fond memories of the mission briefings from *StarCraft*. Characters of command or plot significance – Arcturus Mengsk, The Overmind, or Tassadar, to name a few – would appear as talking portraits, and convey the narrative, context, and exposition required to understand your upcoming task, then would make your objectives clear. Destroy the Zerg hatcheries. Hold this position until reinforcements arrive. Reach the beacon before the Protoss do. The goal was stated cleanly. Nothing else was given to you.

Once you got into the game, you could only see what your units could see, sometimes you'd have a view or marker of where the objective was. Fog of war covered everything beyond your immediate position. You knew where you were starting, and you knew – roughly – what you were trying to accomplish. Everything between those two points was obscured until you went and looked at it yourself.

This is, as it turns out, a precise description of what enterprise delivery actually feels like.

The objective is rarely the hard part. Someone has usually worked out what needs to happen, what the end state looks like – a platform migration, a product launch, an operating model that doesn't exist yet and needs to be built from scratch. The *what* arrives, stated clearly or approximately clearly, and then the briefing ends and you're left standing at the edge of a blacked-out map with your starting units and a waypoint in the distance.

What you do next is the whole job.

The complexity of enterprise delivery is genuinely staggering, and I want to be precise about what I mean by that, because the word gets used casually in ways that undersell it. It isn't just scale – though thousands of users, each with their own workflows and expectations and organizational contexts, are not a small thing. It isn't just the technical stack – though the layers of infrastructure, middleware, integration dependencies, and legacy architecture that most enterprise programs sit on top of constitute their own navigational challenge. It isn't even the process surface area – though the distance from initial concept to deployed, adopted, supported solution is longer and stranger than most people outside it appreciate.

It's all of those things, simultaneously, interacting with each other in ways that don't announce themselves in advance. A dependency you didn't map surfaces three weeks before go-live. A stakeholder who signed off in planning develops reservations when the rollout hits their division. The adoption curve you projected against one set of user behaviors encounters a different set of user behaviors, because users are not the rational actors that adoption models assume they are. The system you built to specification turns out to be not quite the system people needed, because the specification captured what they said, not what they meant.

None of this is solvable with a better Gantt chart.

The orientation that actually works in this environment – and I mean this structurally, not as a personality preference – is the one the Systems essay tried to establish at the start of this volume: zoom out far enough to see the whole map, then navigate it with enough discipline to make forward progress without losing sight of where you're trying to end up. Systems thinking, expressed as operational practice. The fog of war is always going to be there. The question is whether you've sent enough scouts in the right directions before you committed your main force.

This is what full-stack ownership means to me in practice, and it is the capability I believe separates good delivery leadership from excellent delivery leadership.

An average PM can articulate the system they directly inhabit. They know their lane, their dependencies, their stakeholders, their sprint cadence. That's the floor of the job, and it's necessary. A PM who can't navigate their own layer can't navigate anything. But their understanding of their system begins and ends at its borders: they can articulate the inputs and outputs, and the process it runs internally, but they don't understand how their system integrates into a larger whole. That's a problem, because a program at genuine scale doesn't live in one layer. It lives across a stack – from the architectural decisions being made in engineering to the adoption behaviors playing out in end-user populations, with every layer of product ownership, change management, stakeholder alignment, and operational readiness sitting between them. The PM who can only see their own layer is playing with the minimap covered.

The excellent PM – not just the good one – is the one who can move as many levels up or down the stack as the situation demands. That doesn't mean they are going to do everyone's job, but it does mean that they understand it well enough to see where things connect, where they depend on each other, where the fragile joints are, and where the intervention that looks small on the org chart will have the largest leverage on outcomes. The junior PM can articulate their own system. Where a good PM can go a level or two in either direction, the excellent PM goes as far as the problem requires, and knows when they've reached the edge of their useful understanding. That range is what makes the difference between managing a program and leading one.

Which is why the gap between *delivery* and *adoption* deserves its own treatment, because they are not the same problem.

Delivery is the technical accomplishment: the thing exists, it works, it shipped on time and within scope. Adoption is the human one: people are actually using it, in the ways it was designed to be used, with a level of fluency that produces the outcomes the delivery was meant to enable. Conflating these two is one of the most expensive mistakes a program can make. A platform can be technically excellent, delivered on schedule, and still fail at scale because the rollout treated behavior change as a communication problem rather than a systems problem. The fix is never a better training deck. It's understanding the social architecture of the organization well enough to know where trust lives and how to activate it – peer networks, champions, early adopters who can translate the new thing into the language of the people who haven't adopted it yet. That requires a different kind of map than the technical one, and the PM who can only read one of them is going to hit a wall.

This is the orientation I'd carry into game development: not a set of borrowed tools, but a way of reading the problem space before moving through it. The specific stack is different. The stakeholder dynamics have a texture I'll spend real time learning. The release pressures and community feedback loops that define live service production are not identical to what I've navigated before, and I won't pretend they are.

What I have done is navigate ambiguity at scale – repeatedly, in high-stakes environments, without a complete map – and deliver results that held up over time. The specific ambiguity that games delivery introduces is a new layer. But ambiguity at scale is not a new problem for me. It is, specifically, the problem I have spent my career learning to solve.

I'll be up-front about all of this: I have not shipped a raid tier. I haven't built a dungeon. I haven't made a game before. I want to say that plainly, because the reader who has shipped one knows exactly how much

that entails, and I have no interest in minimizing it. The domain knowledge is different. There are things I will need to learn that I do not yet know how to learn, because I don't yet know I need them. That's the nature of every new layer of ambiguity – the unknown unknowns are the ones that matter most, and you can't scout what you don't know to look for.

What I can tell you is what I do when the briefing ends and the map is dark: I send scouts. I map the stack before I commit resources. I find the layers I don't understand and develop enough fluency in them to see where they connect to the ones I do. I treat adoption as a separate problem from delivery, and I don't declare victory until both are solved. And I operate at the level of the people building the thing, not just the thing itself – because the people make the decisions that the thing is built from, and a PM who can't earn their trust can't move the work.

The briefing is always short. The map is always mostly dark.

What matters is what you do when the briefing ends.

PRAXIS

How I Work

Most teams figure out how to work with each other the slow way: months of misaligned expectations, learned friction, the gradual archaeology of discovering what a person actually needs versus what they said they needed in the interview. It doesn't have to work that way. Marshall Goldsmith, in *What Got You Here Won't Get You There*, makes a simple argument: tell people who you are and how you operate, early and directly, before they have to learn it the hard way. I think he's dead on with that idea. So here's mine.

What gets the best out of me

Tell me the mission, not just the task. I operate at a different level when I understand what winning looks like for the organization, not just the project. Connect my work to something that matters and I'll bring more than you asked for. Hollow or purely transactional dynamics will produce a competent version of me. Give me a mission I believe in and you'll get something harder to find.

Give me context up front, then feel free to stand back. I look before I leap, I will read the room, study the stakeholders, and build a mental model before I move. That ramp-up period is not low engagement – it's simply the work before the work. Trust it.

Treat me as a thinking partner, not an executor. I don't just want to carry out decisions. I want to understand them, pressure-test them, and try to improve them. Bring me into the why, not just the what. The best version of this I ever experienced came from a manager early in my career who said: *I don't need to approve your decisions – I've already approved of you.* That's the standard I hold myself to with my own teams. It's also what I respond to most strongly as a direct report.

How I communicate

I think with and through language. Writing is not just how I share thoughts – it's how I structure them. If you want my sharpest thinking, you'll get it in writing. As a result, asynchronous communication is ideal when applicable, not because I can't work live, but because the space to structure my thinking produces better answers than the pressure to produce them in real time.

I will tell you bad news. Not dramatically, not with excessive hedging – but I won't hide a problem once I see it. If I flag something as an issue, take it seriously. I don't cry wolf.

I interrupt sometimes. Not because I don't care – because I do. My brain moves fast and I don't want a thought to vanish before it gets said. If I jump in, it's excitement, not disrespect.

What to watch for

I can over-invest in framing and under-invest in shipping. If something is getting refined past the point of diminishing returns, push me to call it done. The prompt that works best: *what's the one thing we can ship today?*

When the mission feels hollow or priorities go fuzzy, I may disengage unintentionally instead of calling it out. I won't make a scene – I'll just gradually stop generating. A direct "what can I help you move forward?" surfaces it faster than waiting for me to raise it myself.

I'm a systems thinker, which means I'll sometimes keep zooming out to reframe the problem instead of solving it at ground level. If I'm stuck in orbit when the room needs a ground-level answer, pull me back down.

What makes me stay

A mission I genuinely believe in. A manager who extends trust before I've had to prove it. Problems worth solving, not just tasks to execute. Room to grow – not just upward, but deeper into craft. The feeling of belonging to something, not just working somewhere.

I don't need everything to be perfect. I need it to be *aligned*. Give me that, and I'll do more than you hired me for.

The way my brain works

I have ADHD. I'm not disclosing this as a liability to manage – I'm claiming it as the thing that makes me useful.

Here's what ADHD means for me: my brain does not stay in its lane. It is constantly sending out feelers, pulling in information from whatever is adjacent, unexpected, or only loosely related to the problem at hand – and then, without warning, finding out how it connects to something I learned years ago, in a completely different context. That collision, the moment two unrelated things suddenly illuminate each other, is where my best thinking lives.

Look at this book. The epigraph is Tennyson. The BlizzCon essay is structured after Proust.. None of that happens without a brain that is perpetually cross-referencing everything it has ever encountered, looking for the seam where one idea rhymes with another. That lateral connectivity is not incidental to the work. It *is* the work. In a creative organization – one that builds worlds, tells stories, and asks its people to find meaning in the things they make – that kind of mind has genuine value.

The practical corollaries are worth knowing. My focus, when something earns it, is deep and durable. Interruptions in that state are expensive – not because I'm difficult, but because the thread I'm holding is long and getting back to the same place takes time. I forget things that don't get captured immediately, not from carelessness but from the same open-buffer architecture that enables the lateral thinking. Help me capture what matters and I'll return the favor many times over.

Case Study – Aeoncore Personal AI/Services Platform

Overview

Aeoncore is a self-hosted, end-to-end AI and services platform, designed and built from the ground up as a solo project during a professional development period in 2025–2026. The platform delivers local AI inference, private cloud storage, network services, and containerized infrastructure – all running on a single physical node, managed entirely by one person, and operated to production-equivalent standards.

It was not built for commercial deployment. It was built to answer a question: *what does it look like when someone applies real production discipline to a personal project, without a team, without a budget, and without the option to buy their way out of hard problems?*

The answer is Aeoncore.

Context & Motivation

The Catalyst

The project originated from an opportunistic hardware situation. A high-end gaming workstation – Intel i9, 128GB RAM, mixed NVMe/HDD storage, and an NVIDIA RTX 3090 – had become available and was sitting largely idle. The RTX 3090 was the pivotal component. While four years old and two GPU generations behind the current market, its 24GB VRAM capacity occupies a meaningful sweet spot in the modern AI hardware economy: large enough to keep sizable open-weight models resident entirely on-card, which is orders of magnitude faster than offloading into system RAM. A machine that predated the AI era turned out to be well-suited for it.

At roughly the same time, the landscape for open-weight models had crossed a usability threshold. Tools like Ollama had made local model management genuinely accessible. Frameworks like Stable Diffusion had matured. What had previously required enterprise infrastructure – or a cloud vendor's billing account – could now run on well-configured consumer hardware. Building a legitimate local AI platform stopped being theoretical.

The Professional Rationale

Having the hardware removed the capital barrier. What remained was the design and execution challenge, which is where the professional value of the project actually lives.

The build served three simultaneous goals. First, it provided a forcing function for hands-on learning across a technology stack – hypervisors, containerization, networking, AI tooling – that could not be learned as effectively any other way. Reading documentation is not the same as being responsible for an operational system. Second, it produced a tangible artifact: something demonstrable, explainable, and improvable, rather than another theoretical exercise. Third, it delivered genuine utility. The platform had to become part of daily workflows for it to be worth building. A portfolio project that no one actually uses is not evidence of operational discipline.

The Ownership Thesis

Underneath the practical rationale was a more fundamental conviction. When you rely entirely on cloud services, you consume an experience someone else designed. You accept their constraints, their pricing model, their roadmap decisions, and their architectural tradeoffs. Building and operating your own platform forces a different posture – one that requires you to understand the full stack, own the consequences of every decision, and design around real constraints rather than someone else's defaults.

If you want to truly own the experience, you probably need to own the platform.

What began as an experiment in capability became an exercise in ownership.

Design Philosophy

Enterprise Standards at Personal Scale

The guiding principle throughout the design and build was: *enterprise-grade discipline, implemented at personal scale.* This was a deliberate choice, not an aspiration. It shaped every architectural decision, from the initial choice to use a hypervisor rather than a bare-metal install, to the decision to implement monitoring and backups before touching the AI services.

The easiest version of this project would have been a monolithic install: Linux, Docker, Ollama, done. That approach works for a hobbyist. It is brittle, difficult to reason about, hard to scale, and expensive to change. It produces a running system, not an operational platform. The difference matters.

Constraint as Design Variable

Aeoncore was built under a hard constraint: the project budget was effectively zero. No additional hardware purchases. No software licensing. No paid services beyond the electricity bill. Every design decision had to work within what existed.

This constraint was not treated as a limitation so much as a parameter to design within – exactly the kind of parameter that governs real production environments, where budget ceilings are real and the question is never "can we buy our way to a better solution?" but "what is the best architecture given what we have?"

The most significant constraint was power. The existing power supply unit barely supports the RTX 3090 at full load. This means the GPU cannot be joined by a second card without a power supply replacement – a ceiling that caps the compute expansion path and shapes the architectural roadmap accordingly. Acknowledged as technical debt, this constraint is documented and has a known mitigation path: a PSU upgrade when budget allows restores the expansion option. In the meantime, the architecture is designed to maximize the single-GPU configuration, with VRAM allocation managed carefully across workloads.

There is a secondary operational consequence worth noting transparently: because the GPU is allocated exclusively to the Deneb VM at the hypervisor level, local physical management of Deneb during boot requires timing an intervention before GPU passthrough completes – a manual workaround for a constraint that remote management cannot fully address. It is not elegant. It is documented, understood, and stable.

Modularity as Risk Mitigation

Given a $0 budget, the ability to re-architect without re-purchasing was not optional – it was the only available risk mitigation strategy. The platform was therefore designed for modularity from the start: discrete virtual machines for discrete functional roles, containerized services within each VM, and clean interfaces between layers.

This design choice proved its value repeatedly. When the understanding of a subsystem deepened and its configuration needed revision, the change could be scoped to the relevant VM or container without cascading disruption to the rest of the stack. When a networking decision turned out to be wrong, it could be corrected at the network layer without touching the AI services. The modularity was not premature optimization – it was structural resilience against a known pattern of iterative learning.

Architecture

The Physical Node

Aeoncore runs on a single workstation chassis acting as a compact home datacenter:

- CPU: Intel i9 (multi-core)
- RAM: 128GB system memory
- GPU: NVIDIA RTX 3090 (24GB VRAM)
- Storage: Mixed NVMe (primary) and HDD (bulk/NAS)
- Network: Internal LAN with zero public-facing exposure

The Hypervisor Layer

Proxmox VE was selected as the hypervisor for several reasons: open-source licensing (no cost), strong community documentation, native Linux compatibility, built-in LXC container support alongside full VM management, and the ability to support multi-node clustering if a second physical machine is added in the future.

The decision to run a hypervisor rather than a direct install was deliberate. Service isolation at the VM level allows resource allocation to be adjusted, workloads to be migrated between hosts, and individual VMs to be snapshotted and restored independently. It is the difference between a house built to code and one that was not – the former can be added to; the latter has to be torn down.

The Functional Zones

The platform is partitioned into three specialized virtual machines, each with a defined role:

Sirius is the general services VM, running on Ubuntu LTS. It handles network infrastructure (Tailscale mesh VPN, Nginx Proxy Manager, DNS, Pi-hole ad blocking), platform monitoring (Beszel, Uptime Kuma), container management (Dockge, Dozzle), backup orchestration (Kopia), and user-facing utility services (a unified homepage portal, Linkwarden).

Vega is the storage VM, running TrueNAS SCALE with direct hardware passthrough of the physical HDDs. This enables a proper ZFS mirror array for data integrity, provides a shared network storage location accessible by the platform's other services, and operates as a private cloud for user devices – with sync clients on desktop and mobile – without a single byte of data leaving the local network.

Deneb is the AI compute VM, running a base Linux environment with exclusive GPU passthrough of the RTX 3090. It carries twice the CPU allocation and six times the RAM of its peer VMs, and runs the full AI service stack: Ollama and Open WebUI for text inference (the "Tau" front-end), and Stable Diffusion with ComfyUI for image generation (the "Ceti" front-end).

Networking & Security

All user-facing services are internal only – there are no public-facing endpoints. Remote access is handled exclusively through Tailscale, a zero-trust mesh VPN that authenticates at the device level. Access to the network IS authentication; if a device can route to the platform, it has already been verified. This approach was chosen deliberately to minimize attack surface – not from a position of networking expertise, but from a position of honest risk assessment. The platform is most secure when the default posture is closed.

Containers & Orchestration

Services within each VM run as Docker containers managed via Docker Compose. Kubernetes was evaluated and rejected: for a single-node configuration, it adds architectural complexity with no immediate operational return. Docker Compose provides rapid deployment, clean versioning of service configurations, easy rollbacks through image pinning, and the ability to spin up new services in minutes rather than hours.

Build Execution

Sequencing the Build

The project was built in deliberate phases, each validating a layer before the next was introduced:

The first phase established the bare-metal foundation: hardware assembly, Proxmox installation, initial VM configuration. Nothing ran on it yet. The foundation had to be stable before anything was built on top of it.

The second phase stood up Sirius and the general service stack – monitoring, networking, container infrastructure, backups. The explicit decision was made to get observability and backup systems operational *before* building any user-facing services. This is production thinking: you do not build on a foundation you cannot see and cannot recover from.

The third phase added storage. Hard drives were added to the chassis, Vega was provisioned with TrueNAS, ZFS was configured, and the private cloud sync was established. Shared storage became available to the rest of the platform.

The fourth phase provisioned Deneb – allocating the GPU, building out the AI service stack, and bringing Tau and Ceti online. By the time AI services were being configured, the platform beneath them was already stable, observable, and backed up.

The Reality of Iterative Learning

The build process was iterative, reflective of what happens when a practitioner is learning a technology stack while simultaneously architecting and operating it.

The overhead cost of this approach was real. Planning and principal build occurred before agentic AI tooling had reached the capability level where it could maintain coherent architectural context across a multi-layer system build. That context had to be maintained at the human level: in documentation, in design notes, in a mental model that had to stay current with each change to the underlying stack. It was a meaningful cognitive load, and it produced a genuine appreciation for what that kind of tooling – when it matures – will change about solo infrastructure work.

The Docker Compose configurations were the most frequent site of iteration. Each new service layer had to be integrated with the existing compose files, and each integration surfaced new dependencies or conflicts that required adjustment. The Deneb AI stack was particularly involved: Ollama, Open WebUI, and the Stable Diffusion layer needed to see each other and share GPU resources without over-allocating VRAM. Configuring SDXL to release the GPU when idle – rather than holding VRAM permanently – required understanding the resource lifecycle at a level the initial configuration had not accounted for. The models, once configured correctly, coexist cleanly. Getting there required several iterations.

The backup configuration provided the clearest example of the find-and-fix pattern. The initial setup backed up VM data to the NAS – sensible in concept. What it actually did was attempt to back up the NAS's own operating data to itself: a circular configuration that produced no useful backups and, worse, created a false sense of coverage. The failure was not immediately visible because the process completed without obvious errors. It surfaced during a configuration audit. The correct approach – backing up VM snapshots through Proxmox directly, with NAS data handled separately through ZFS native snapshots – was implemented, tested, and has been stable for months. The lesson was not that mistakes happen; it was that backup configurations require validation, not just completion. A backup you have not tested is not a backup.

Other iterations were less dramatic but equally characteristic: CPU and RAM reallocations between VMs as actual utilization patterns became clear, IP addressing and DNS namespace adjustments as the internal service map grew, authentication and permissions tuning as the user base expanded beyond a single person.

In a production environment with SLA obligations, this rate of change would require change management processes to control risk. Aeoncore's operating context – no SLA guarantees, no external dependencies – allowed the iteration to be aggressive. The platform was designed to tolerate it. That is not an excuse for the changes; it is the reason the modular architecture was worth building in the first place.

Current State & Operational Status

What 1.0 Delivered

Aeoncore 1.0 is operational. The core platform services – monitoring, networking, backup, storage – run without active management. The NAS and private cloud sync are used daily. The AI services are available on demand for research, experimentation, and workflow development.

The AI services are not in heavy daily use. This is worth stating directly, because the honest accounting of what a personal AI platform delivers matters more than the theoretical accounting. What the platform provides is *readiness*: the ability to run a local model, test a new capability, prototype an agentic workflow, or evaluate a new tool without spinning up a cloud account, writing a check, or waiting for provisioning. The marginal cost of experimentation is near zero. That changes what experimentation is possible.

The 2.0 Roadmap

The 2.0 roadmap is documented and partially in progress, with execution gated primarily on budget constraints that remain unresolved. Planned features include the Tau/Ceti unified pipeline (natural language image generation without leaving the chat interface), autonomous web search with citation (moving from static inference to a dynamic research partner), Obsidian knowledge base integration via local RAG (turning private notes into a searchable AI-backed resource), and a dedicated Proxmox Backup Server node for off-primary resilience.

Several of these features require modest capital expenditure – hardware for the backup node, specifically – that is deferred until the professional employment situation resolves. Program-level planning has continued in the interim. The platform is designed to receive these capabilities without re-architecting the core; the 1.0 architecture was built with 2.0 in mind.

What This Demonstrates

Production Discipline Transfers

The argument Aeoncore makes is not that building a home server is impressive. Plenty of people build home servers. The argument is that the discipline applied here – hypervisor-first architecture, service isolation, observability before services, modular design, documented iteration – is the same discipline that governs production infrastructure decisions, because it *is* production infrastructure discipline, applied at a different scale.

The scale is personal. The standards are not.

A platform built this way can be explained to an engineer, evaluated by an architect, and operated by a practitioner. It was designed with all three audiences in mind, because the person building it was wearing all three hats simultaneously.

Full-Stack Ownership

Program and product management roles frequently create distance from the implementation layer. That distance has costs. A program manager who cannot read an architecture diagram cannot evaluate whether the architecture serves the program. A product manager who has never operated a production service does not fully understand what "operational overhead" means when an engineering team raises it as a concern.

Aeoncore closed that distance deliberately. Every decision in the stack – from hypervisor selection to VRAM allocation strategy to backup configuration validation – was made, owned, and corrected by the same person. There is no part of this platform that someone else built and this author simply managed.

That is a professional posture, not a technical credential. The specific technologies will change. The posture does not.

The Platform as Capability

The final argument is forward-looking. Aeoncore exists not just as a completed project but as a standing capability – a platform that makes future experimentation fast and future proof-of-concept work low-friction. New models can be evaluated in hours. New agentic workflows can be prototyped against real infrastructure. New integrations can be tested without cloud accounts or API costs.

The value of building the platform is not only what it has already done. It is what it makes possible to try next.

Technical Stack Summary

Layer	Technology	Role
Hypervisor	Proxmox VE	Physical resource management, VM isolation
General Services VM	Ubuntu LTS (Sirius)	Networking, monitoring, containers, backups
Storage VM	TrueNAS SCALE (Vega)	ZFS array, shared storage, private cloud sync
AI Compute VM	Linux + RTX 3090 (Deneb)	LLM inference, image generation
Text Inference	Ollama + Open WebUI (Tau)	Local LLM front-end
Image Generation	Stable Diffusion + ComfyUI (Ceti)	Local image generation
Networking	Tailscale	Zero-trust mesh VPN, device authentication
Reverse Proxy	Nginx Proxy Manager	Internal service routing
Container Mgmt	Docker Compose + Dockge	Service deployment and management
Monitoring	Beszel + Uptime Kuma	Resource and uptime observability
Backups	Proxmox snapshots + Kopia + ZFS	Multi-layer data resilience

---

Aeoncore is a living platform. This document reflects its state as of May 2026.

Case Study – BP M365 and Digital Transformation

Enterprise Collaboration Transformation – Change Management & Adoption Strategy

---

Overview

In the summer of 2018, BP engaged Prosource to provide program leadership for a strategic initiative: converting a technically-deployed Microsoft 365 platform into something employees actually used. The engagement spanned six months and covered a global workforce of approximately 100,000 users across corporate employees and contractors.

The core objective was specific and measurable: move Teams adoption from 25% of users engaging weekly to 65% of users engaging daily. That is not a modest target. It is a 2.5x expansion of the user base and a 5x increase in the frequency standard – applied simultaneously, across one of the largest corporate digital footprints in the energy sector.

The program was delivered on objective.

Context

The Situation

BP had completed a technically successful migration of their mail and productivity infrastructure into Microsoft 365[36]. The licenses were active, the infrastructure was running, and from an IT perspective the project was closed. Employees opened Outlook and saw Outlook. Nothing had visibly changed.

That was the problem.

M365's value is not in its email client, it is in the connective tissue – Teams for communication and collaboration, Planner for task management, Forms, To Do, Yammer, and for specific populations, Project Online and Visio. BP had purchased E5 licensing across the enterprise, which included all of it. Almost none of it was being used.

Teams was the strategic priority. Microsoft had announced the deprecation roadmap for Skype for Business, and BP needed their workforce on Teams before that deadline arrived. The adoption numbers at engagement start were clear and concerning: roughly 25% of the 100,000-user population was opening Teams at least once per week. The target was 65% daily – a threshold that represented genuine behavioral change at scale, not surface-level awareness.

Organizational Complexity

The program's organizational structure was itself a constraint worth documenting. BP maintained two internal groups with relevant ownership: the Collaboration team, responsible for M365 engineering, support, and day-to-day platform operations; and the Digital Workplace team, running a broader enterprise transformation program from London headquarters. These two groups had overlapping mandates and a documented tension between them – the Head of Collaboration had made clear from the outset that she believed this program belonged to her team, while Digital Workplace held nominal authority over the initiative.

The role Prosource was hired to fill sat directly between them. Operationally reporting to Digital Workplace, doing most of the hands-on work in partnership with Collaboration, and navigating the friction

[36] At the time, Microsoft 365 was called Office 365, as the rebranding occurred in 2020. I have kept the modern name to alleviate potential confusion.

between them without the organizational authority to resolve it. This was not incidental to the job. It was the job.

Approach

The before state was straightforward to diagnose. The platform had been deployed with baseline training documentation that employees largely ignored, minimal internal marketing, and no behavioral change strategy. The technical implementation was complete. The adoption work had not begun.

The program required building that work from scratch – design, execution, and measurement – across a global, culturally diverse, time-zone-distributed workforce.

Adoption Strategy Design

The first task was to replace the existing partial draft of a strategy – described by a Digital Workplace colleague as "rather dry" – with something that could actually drive behavior at scale. The revised strategy organized around several parallel tracks, running simultaneously rather than sequentially.

The diagnosis underneath the strategy was that this was not a technology problem. The platform worked. The challenge was behavioral: employees had existing habits, existing tools, and no compelling reason to change. Any adoption program that led with features rather than workflows was going to fail the same way the original documentation had failed. The approach had to lead with people.

Internal Marketing and Communications

Working directly with BP's Corporate Communications function, a sustained communications campaign was developed and executed across company-wide channels – email newsletters, internal announcements, and targeted messaging to specific populations. The goal was not to announce the product. It was to make the product visible as something relevant to how people actually worked, which required translating platform capabilities into recognizable workflow language.

Workflow Workshops

A series of targeted workshops was designed and run with cross-functional teams drawn from across the enterprise. The design was deliberate: each workshop began not with a product demonstration but with a workflow problem. Teams would surface their actual coordination and communication pain points, and the workshop would work through how M365 tools addressed them specifically.

The recurring diagnosis was consistent across teams: internal communication and coordination were harder than they needed to be, and employees weren't sure how their tools were supposed to help. The workshops produced small case studies – real teams, real problems, real solutions – that could then be redistributed across the enterprise as proof points. Adoption at scale is easier when employees can see peers from recognizable teams describing recognizable problems being solved.

Champion Network

A network of internal advocates – named Digital Pioneers – was established across the enterprise. Candidates were identified through a combination of existing intelligence from a prior Prosource program and warm introductions from a colleague. The selection criteria prioritized employees who were already technically engaged and organizationally influential: people whose adoption would be visible to their teams and whose enthusiasm would be credible.

The Champions served as a distributed change agent network – receiving direct communication and support, then carrying the program's message into their own organizations. This model extended the program's reach far beyond what a central team could achieve through direct contact alone.

Flagship Events

Two flagship adoption events were designed and executed – one at BP's US headquarters in Houston, one at London HQ. The format was intentionally more substantive than a standard lunch-and-learn: a half-day structure with a keynote and breakout sessions, designed to give employees a meaningful engagement with the platform rather than a surface-level awareness exercise. The Houston event was run directly. The London event was led by a Digital Workplace colleague.

Outcomes

The program delivered against its primary objective. Teams adoption moved from approximately 25% of users engaging weekly to 65% of users engaging daily – a change in both the population reached and the behavioral standard measured. Across a 100,000-user footprint, that represents roughly 40,000 additional people integrating a daily-use collaboration platform into their working lives.

Secondary platform tools – Planner, Forms, Yammer, and others – saw corresponding lift as the adoption program created awareness across the broader M365 suite.

At the six-month mark, the Collaboration team assessed that they had sufficient internal capability to continue the program independently. The handoff was clean. The work continued under their ownership.

A senior member of the Digital Workplace leadership offered the following assessment:

"He is very strong on change management, its challenges and its key concepts, and he was full of enthusiasm for the tasks set. He quickly grasped the scope and scale of the task at hand during our early meetings, and identified many areas that hadn't yet been fully thought through. He took a rather 'dry' existing partial draft of the strategy, and really brought it to life in both its written form and in how he articulated it in presentations. He gets on well with his fellow team members and with stakeholders. He is a solid and technically savvy change manager."

What This Engagement Demonstrated

Ambiguity Resolution at Scale

The program arrived with a clear metric and an unclear path. No playbook existed for this specific combination of platform, organization, and behavioral target. The work required rapidly synthesizing an

understanding of the organization's structure, its political landscape, and its employees' actual working habits – then designing an approach that could operate across all of it simultaneously.

The ability to move from ambiguity to clarity to plan to execution, without waiting for the ambiguity to resolve itself, was the core competency the engagement demanded.

Behavior Change Is a People Problem

The framing that shaped everything else: this was not a technology program. The technology was already deployed. The program's actual subject matter was human behavior – specifically, the gap between what a tool can do and what people will choose to do with it.

Closing that gap required understanding employees not as users to be trained but as people with existing habits, legitimate workflow concerns, and limited patience for solutions that did not visibly address their problems. The workshop model worked because it started there. The champion network worked because it leveraged existing social trust rather than trying to manufacture it from outside. The communications worked because they spoke in workflow language rather than product language.

Technology adoption at enterprise scale is a product problem wearing a change management costume. The costume matters less than the underlying diagnosis.

Navigation Without Authority

The organizational structure placed this role in a position of significant influence with limited formal authority – caught between two internally competing teams, neither of which fully owned the space the program needed to operate in. Navigating that required an ongoing calibration: enough alignment with each team to maintain cooperation, enough independence to keep the program moving when the teams' priorities diverged.

This is not an unusual position for a senior program manager. It is, in fact, the standard one. The ability to operate effectively inside organizational friction – without requiring it to be resolved before the work can proceed – is a prerequisite for any complex program environment.

Scale Changes the Work

One thing this engagement made clear that smaller programs could not: scale is not simply more of the same. At 100,000 users, every decision about communications, targeting, and sequencing has downstream effects that a 5,000-user program would never surface. The same approach that works in a mid-sized organization fails at enterprise scale – not because it is wrong in principle, but because the system it is operating in behaves differently.

Understanding how scale changes a program's dynamics – and designing for it explicitly rather than discovering it reactively – is a capability this engagement developed in ways that no smaller program could have.

Engagement period: Summer–Winter 2018. Client: BP (via Prosource).

Case Study – CVS Enterprise PC Refresh

Overview

In February 2022, CVS Health engaged an outside program manager to lead their annual PC refresh cycle – a recurring enterprise initiative to replace end-of-life hardware across all three business units: CVS Retail, Aetna Healthcare, and Caremark Prescription Benefit Management. The program ran through August 2022.

The delivery target was 18,000 devices. The deadline was end of year.

The program closed at 106% of target in July – five months ahead of schedule.

Context

The Program

CVS runs a PC refresh cycle every year. The eligible device population is determined annually against a set of criteria – hardware age, installed RAM, CPU generation, operating system status – that typically surfaces roughly 20% of the enterprise device footprint for replacement. Those devices are catalogued, a target is set at a practical 100% threshold, and a program is stood up to execute against it.

The 2022 cycle targeted 18,000 devices across the full enterprise, with approximately 20,000 available as a stretch goal. Six projects ran in parallel: a Refresh project and a Growth project (covering organic headcount expansion) for each of the three business units, each with its own IT organization, its own desktop support structure, and its own operational context.

The Structural Challenge

CVS runs this program every year. It almost never retains the same program manager twice.

The downstream consequence is a standing knowledge debt. Each incoming program manager begins from approximately zero – rediscovering undocumented institutional knowledge, re-establishing relationships with the same teams, and rebuilding a program framework that should have been preserved and improved year over year. It is an organizational inefficiency that compounds the startup cost of every cycle.

That context is worth naming because it shaped the initial weeks of the engagement: a significant portion of early effort went into knowledge acquisition that a mature program management practice would have made unnecessary. The program succeeded despite this pattern, not because of it.

Program Structure

The Team

The program operated across three business unit IT organizations, each structured with a desktop manager, area leads across major office locations, and distributed desktop technician teams. In total, approximately 100 people touched the program across its lifecycle – though not simultaneously.

Execution was deliberately staged. Rather than activating the full team at once, cities and office locations were sequenced: a cohort would be activated, the refresh work completed, and that cohort would stand down before the next was brought online. At any given time, the active working group consisted of the three

business unit desktop managers, approximately five area leads, and twenty technicians. This approach kept coordination overhead manageable and allowed focus and accountability to concentrate on the active geography.

The desktop managers were the primary day-to-day interface. Daily standups with that group kept deployment on track and surfaced issues before they became schedule risks. Weekly team-wide meetings ensured leads and active technicians maintained situational awareness across the program.

Procurement and Logistics

The procurement pipeline – working with Dell's account team and CVS's internal procurement lead – was already established when the program began. The program management role was to oversee it: ensuring devices were ordered to match the deployment schedule, shipped to the correct locations at the correct times, and available for technicians when each geographic cohort activated. Timing mismatches between procurement and deployment would have created idle time or schedule compression; neither occurred.

Imaging and local deployment were handled by the business unit technicians and their leads, who carried the operational knowledge of their end users and locations. The program manager role was coordination and sequencing, not hands-on technical execution.

Execution

Cost Accounting and Executive Reporting

The most demanding ongoing responsibility was financial – specifically, resource utilization forecasting and program cost accounting across six parallel projects. Weekly status reports went to VP-level leadership, one level below the CTO. A monthly roundtable required a formal presentation summarizing program health, deployment progress, financial actuals versus forecast, and forward projections.

Keeping six parallel workstreams financially visible and accurately forecast – while deployment geography shifted week to week – required a discipline around the numbers that did not allow for estimation lag. Executives making resource decisions needed accurate data, not approximations. That standard was maintained throughout.

Program Documentation and Visibility

The full documentation suite – RAID log, stakeholder register, deployment dashboards – was maintained as living documents rather than static snapshots. The intent was to keep information actionable: leadership could pull current status at any point without waiting for a scheduled update, and the team leads had a shared source of truth for deployment priorities and sequencing decisions.

Parallel Workstream Coordination

Six projects running in parallel across three organizational structures introduced coordination surface that a single-project program does not have. Dependencies between the Refresh and Growth projects within each business unit required ongoing alignment – a device procured for a Growth deployment could not be

reallocated to a Refresh without cascading effects on both project plans. Maintaining that discipline across the full program, at the pace the delivery schedule required, was the primary coordination challenge.

Adjacent Work

Beyond the core refresh, the program contributed to a parallel initiative: updating CVS's device decommission process, including automating the removal of retired machines from Active Directory. This was not the primary program objective, but proximity to the refresh cycle made the collaboration natural, and the process improvement had downstream value for future cycles.

Outcomes

The 18,000-device target – the defined 100% threshold – was reached in July 2022, five months ahead of the December deadline. By contract end in August, the program had deployed 106% of the primary target and contributed to stretch goal progress across the business unit Growth projects.

The program closed on budget. No devices were delayed past their scheduled activation window. Executive reporting was current and accurate throughout.

The contract was not renewed – because the work was done.

What This Engagement Demonstrated

This was not a complex program. It is worth saying that directly, because inflating a clean execution into something it was not would undermine the credibility of the case studies that deserve more analytical weight.

What CVS demonstrated is different, and in its own way equally important: when the job is delivery, the delivery happens. On time, on budget, at quality, at scale – without drama, without heroics, and without the kind of fire-fighting that indicates a program being managed reactively rather than proactively.

Large programs run smoothly when they are structured correctly from the start. The staging model kept coordination overhead manageable. The financial discipline kept leadership confident and informed. The documentation kept the team aligned without requiring constant escalation. None of these are complicated ideas. Executing all of them simultaneously, across six parallel projects and a distributed team of approximately 100 people, for eight months, is the demonstration.

The program also surfaced a genuine institutional recommendation that CVS did not act on: retaining program continuity year over year. The startup cost of a rotating program manager is real, it is measurable in early-cycle inefficiency, and it is entirely avoidable. That observation was made. What CVS chose to do with it was their decision.

Engagement period: February–August 2022. Client: CVS Health (via contracting agency).

Curriculum Vitae

Donald L. Coles
dlcoles@gmail.com · 713-503-6456
linkedin.com/in/donaldcoles · aeon.ghost.io
The preceding pages make the case. This is the official record.

Experience

Platform Architect & Product Owner · Aeoncore

September 2025 – Present · Independent Project

- Designed and built a self-hosted AI and services platform end-to-end: architecture, backlog, iterative delivery, operations, and ongoing maintenance – solo, no budget, production standards throughout
- Full-stack hands-on fluency across virtualization, containerization, AI/LLM integration, networking, and security – applied to a system in daily active use
- Built for real reliability: automated backups, monitoring, redundancy, and upgrade paths equivalent to production-grade infrastructure

Delivery Program Manager · Sonic Automotive

July 2024 – November 2024 · Remote · Contract

- Inherited a multi-site IT refresh in active distress across 100+ dealerships; diagnosed failure points across six regional leads and 50+ field technicians and rebuilt delivery around clear ownership, sprint cadence, and formal validation checkpoints
- Designed SMS-first workflow for frontline technicians operating in low-tooling environments, replacing an inconsistent verbal process that was the primary source of delays and rework
- Reduced escalations 50%+; closed full program scope a full month ahead of revised schedule

Senior Program Manager · United Airlines

July 2023 – November 2023 · Houston / Remote · Contract

- Enterprise program lead across infrastructure and facilities modernization; bridged operational, technical, and executive stakeholders across distributed initiatives
- Directed discovery and early planning for a 500,000 sq. ft. aircraft parts distribution hub; took over and stabilized a major hangar retrofit at PDX mid-execution
- Partnered with analytics teams on TSA compliance reporting at Newark; developed custom PMO toolset providing visibility across disparate workstreams

Senior Program Manager · Sysco

September 2022 – April 2023 · Houston / Remote · Contract

- Led program management for Omni, Sysco's strategic initiative to realign restaurant customers across its regional OpCo network – a multi-year migration affecting hundreds of thousands of customer relationships
- Co-designed the multi-year migration roadmap with regional leadership: wave sequencing, dependency management, milestone governance, and quality checkpoints across the full delivery lifecycle
- Facilitated cross-functional discovery with stakeholders across operations, finance, sales, HR, and technology; built Smartsheet-based reporting connecting regional leadership to corporate PMO governance

Technical Program Manager · CVS Health

February 2022 – August 2022 · Remote · Contract

- Led enterprise endpoint modernization across CVS Retail, Aetna, and PBM: three distinct organizations, 18,000+ devices, nationwide – 100% on schedule and under budget
- Synchronized delivery across all three business units, managing separate change management requirements and technical constraints under unified program governance
- Applied hybrid delivery: Agile sprint cadence for execution and prioritization; waterfall milestone reporting for executive, compliance, and audit stakeholders

Program Manager, Digital Transformation · BP

August 2018 – March 2020 · Houston

- Directed global M365 rollout for 100,000 users; designed the Champion Network, training, communications, and governance frameworks that increased Teams adoption from 25% to 66% within six months
- Built adoption metrics and feedback loops connecting real-time engagement data to rollout decisions, enabling evidence-based program management across worldwide regions
- Scope expanded based on demonstrated performance: transitioned into infrastructure readiness and data remediation leadership for multi-billion-dollar M&A divestitures, coordinating technical, legal, and operations stakeholders under active compliance constraints

Cloud Practice Manager · InterDyn BMI

February 2017 – December 2017 · Houston

- Led practice operations and delivery management for a Microsoft Dynamics partner during a leadership transition; owned full practice P&L across 30–40 concurrent client projects
- Delivered 60+ Azure-hosted Dynamics implementations within nine months; maintained delivery accountability, quality standards, and team morale through organizational transition
- Served as delivery authority in client-facing and pre-close engagements; provided vCIO advisory services on cloud strategy, migration, and cost optimization

Education & Certification

University of Houston-Downtown • Bachelor of Science, Interdisciplinary Studies (2013) Project Management Professional (PMP) • Project Management Institute (2016 – current) Six Sigma Yellow Belt / Lean • International Association for Six Sigma Certification

---

Core Competencies

Delivery & Program Management Agile • Scrum • Kanban • Hybrid Delivery • SDLC • Release Management • Sprint Planning • Backlog Management • Roadmapping • Dependency Management • Risk Management • Stakeholder Management • Change Management • PMO Governance

Product & Systems Thinking Platform Thinking • Product Discovery • Requirements Definition • Iterative Delivery • Metrics-Driven Decision Making • User-Centric Design • Technical Tradeoff Analysis • Architecture Planning

Technical Fluency Python • JavaScript • HTML/CSS • SQL • Linux • Docker • Kubernetes • Git • Azure • Microsoft 365 • SharePoint • Proxmox • AI/LLM Platforms • Self-Hosted Infrastructure

Tools Jira • Confluence • Notion • Smartsheet • Microsoft Project • Visio • Teams • Excel

End Matter

Afterword – Moonshots

Moonshots, by definition, don't have good odds. I am under no illusions about this one. The standard recruiting pipeline exists for a reason, and the person most likely to receive this book has a full calendar and limited patience for unsolicited materials, however carefully bound. I know that. I made it anyway – because the alternative is deciding that something matters to you, running the probability calculation, and concluding the numbers don't warrant the attempt. I wasn't okay with giving up before I'd started.

So the question was: "what would it look like to actually try?"

Not to apply through Workday, or to submit a resume through appropriate channels and hope to get a response. To really try – to make my case the way it deserved to be made, in a form that couldn't be automated, parsed, reduced to a keyword count, filtered, or ignored. To do the thing that the standard process makes very difficult: let the work speak first, and trust that the right person would hear it.

That's what you've been reading.

What it's been trying to say – and what I'll say plainly now, since the evidence is already behind us – is this: I am a program manager and delivery leader who has spent his career building, running, and rescuing large-scale technology programs, and I have spent thirty years loving the worlds that Blizzard builds with the specific, informed love of someone who gets it. I am not looking for a job that happens to have games in it. I am looking to bring what I know how to do to the place that makes the things I care about most.

I want to work at Blizzard. Building and delivering the things that make these worlds real – managing the complexity, holding the timeline, keeping the people moving in the same direction – so that the people whose genius creates them can do that without the machinery breaking underneath them. That is the work I know how to do. That is the work I want to do here.

I watched Chris Metzen walk back onto a stage in Anaheim in 2023, and something that had been suppressed for a long time got loud again. Not because of who he is, though that matters. Because of what he was: a man who had found his way back to the work that meant something, who was lit from within by it, who couldn't have hidden that even if he'd tried. I recognized it. I used to have it. I've spent the time since building my way back toward it. This book is the proof of that direction. And a request for the conversation it warrants.

Thank you for reading it.

Putting the War in *World of Warcraft*[37]

On Emotional Design, Moral Weight, and the Raids Nobody Puts at Number One

There is a list that exists in approximately ten thousand variations. It's found on gaming websites, forum threads, YouTube videos, and Reddit comment sections. It is the list of the best raids in *World of Warcraft*. And while the specific order shifts depending on who is writing and what era they raided in, the top tier is remarkably stable. Ulduar. Icecrown Citadel. Throne of Thunder. Black Temple. These are the canonical answers, the ones that win by consensus, the raids that defined what raiding could be.

I want to talk about two raids that are almost never on that list, or when they are, sit near the bottom of it. I want to argue that they might be doing something more sophisticated than any raid above them – and that the reason they don't rank higher reveals something important about how we evaluate designed experiences, and how badly our vocabulary is failing us.

The raids are the Siege of Orgrimmar, the climax of the *Mists of Pandaria* expansion, and the Battle of Dazar'alor, from *Battle for Azeroth*. Neither is widely celebrated. Neither would most players cite as their favorite. And yet both have lodged themselves in my memory in a way that Ulduar, for all its mechanical brilliance, never quite has.

When players discuss what makes a raid great, the conversation almost always lands in the same places. Boss design. Encounter variety. The elegance of the hard mode structure. How satisfying the progression felt. Whether the trash was interesting or merely a tax on your time between the real content.

These are legitimate criteria. They are also almost exclusively mechanical criteria. And Ulduar earns its throne by mechanical standards that remain genuinely breathtaking. The hard modes in Ulduar were invisible – you triggered them by doing something in the fight itself, not by selecting a difficulty from a menu. The bosses were creative. The pacing was nearly perfect. It is, on those axes, a near-flawless piece of design.

But here is the thing about Ulduar that nobody puts in the tier list justification: it never asks you to feel anything about Ulduar.

Yogg-Saron is a spectacular boss encounter. The Old God corruption mechanic, the descent into madness, the final chamber – these are extraordinary pieces of fight design. But you never knew Ulduar before it was corrupted. You have no before-picture. You arrive, you discover something terrible has happened, you progress toward stopping it. The location is magnificent, the stakes are cosmically high, and at no point does the game reach into your accumulated history with the world and pull something out that you didn't know was there.

The raids I want to talk about do exactly that. Constantly. Mercilessly.

Most raids position you as a small strike force. However enormous the stakes – a sleeping Old God, a demon lord, a death god sitting on a throne of ice – the framing is intimate. You are the chosen few who can go where armies cannot. You slip through defenses. You delve into depths. You are, in the most literal sense of the word, adventurers.

[37] This essay is an original composition for this manuscript.

The Siege of Orgrimmar is something different from the moment you zone in.

The portal into the instance is not a mystical gateway in some neutral location. It is a hole ripped open in the Vale of Eternal Blossoms – a zone that players spent the entirety of the Mists of Pandaria expansion learning to love, a place of extraordinary stillness and beauty that existed as a kind of sanctuary at the center of the continent. And it has been destroyed. The sha corruption spreads from the wound in the earth like an infection. The Vale that was pristine is now a scar, and the entrance to the raid is the scar itself.

This is not an accident. Before you have fought a single boss, before you have cleared a single trash pack, the game has already asked something of you. It has already reached into your accumulated experience of this expansion and twisted it.

What follows is a progression unlike most raids. The first four bosses take place in the mines beneath the Vale – you are descending into the source of the corruption, fighting your way deeper. Then an archmage teleports you to the outskirts of Orgrimmar, and the next two bosses are about one thing: breaching the front gates of the Horde capital. You are not sneaking in. You are not finding a side entrance. You are hammering at the front door of the enemy's most recognizable city, and when it opens, you pour through.

You are the tip of a spear. You can feel the weight of everything behind you.

This is the feeling I keep coming back to when I think about what separates these raids from the ones above them on every list. It is a feeling of motion. Of military advance. Of momentum that has its own gravity. You encounter resistance in the form of a boss fight, you overcome it, and then you continue your advance. There is an implied army at your back. There is an implied race against time at your front. Garrosh Hellscream is right now, in this moment, consolidating power – corrupting his city further, preparing something terrible, buying time with the soldiers he sends against you. He has agency. He is not waiting for you to arrive. Every boss in the middle of the raid is buying him time, and you can feel that.

For a game called *Warcraft*, it is remarkable how rarely you feel like you are at war.

These raids are the exception.

There is a concept in psychology about the human brain's extraordinary sensitivity to change. We are difference-spotters by nature, wired to notice when something in a familiar environment has shifted. It is a survival instinct, and it runs deep.

Game designers rarely get to use it. Building familiarity takes time – sometimes years. Most games can't afford that investment. But *World of Warcraft* has been running for two decades, and players carry enormous accumulated experience with its world. The Siege of Orgrimmar spends that accumulated experience like currency.

The front half of the raid takes place in locations players already know. The corrupted Vale. The gates of Orgrimmar. The main thoroughfare of the city itself – the place you've landed on a flight path a hundred times, the auction house, the bank, the familiar layout of enemy territory that you've navigated since the beginning of your relationship with this game. And in every location, something is wrong. The architecture is the same and everything else is different. The sha corruption bleeds through the streets. Garrosh's loyalists have remade the city in his image. The Orgrimmar you knew is gone, and what replaced it is a monument to what happens when a faction turns on its own values.

Your brain does the comparison automatically. You don't decide to feel the wrongness – you just feel it. The before-picture exists in your memory, and the after-picture is what you're standing in, and the distance between them is the emotional engine of the entire raid.

The Battle of Dazar'alor operates on the same principle, but the city in question is one Alliance players spent a single expansion exploring rather than years. The designers compensated by making the city itself extraordinary – the golden pyramid of Dazar'alor, tiered and magnificent, is one of the most visually striking locations in the game's history. You climb it. You advance through its levels. You fight your way into the throne room at its heart. The geography is the narrative: progress is conquest, and the city communicates its grandeur and its age with every step you take through it.

Both raids understand something that dungeon-crawl design usually misses: a location is not just a backdrop. It is a character. And like any character, it can be wounded, transformed, lost. The best raid design makes you feel that loss.

Here is a complication that neither raid resolves, and that I think is part of why they resist easy ranking.

Both raids are experienced differently depending on which faction you play.

In the Siege of Orgrimmar, the Alliance player is a conqueror. You are the invading force, and Orgrimmar is enemy territory that has fallen to something monstrous. You are delivering justice. The moral position is clear, the momentum is righteous, and when you finally reach Garrosh, you are the instrument of accountability descending on a tyrant. It is a satisfying story.

The Horde player is in a civil war.

The people you're fighting through are your people – or were, before they chose Garrosh's side. The city you're tearing apart is your city. The NPCs around you aren't obstacles; they're your faction's loyalists who made the wrong choice. You are not liberating Orgrimmar. You are dismantling it from the inside, hoping there will be something worth saving when you're done.

Same bosses. Same hallways. Completely different emotional story.

The Battle of Dazar'alor takes this further by mechanically enforcing the perspective shift. The raid is structured so that both factions play a portion of the content as members of the opposing faction – not just thematically, but literally, your character model changing, your racials changing, your place in the story changing. For three bosses in the middle of the raid, you are playing as the enemy.

As the Alliance, you breach the palace and reach King Rastakhan's throne room. Genn Greymane demands his surrender on behalf of Anduin Wrynn. And Rastakhan – who has ruled the Zandalari Empire for over two hundred years, who watched his patron loa die and made a desperate bargain with the loa of death to keep his kingdom alive, who has spent an entire expansion fighting to protect his people – refuses. He will not bow. Zandalar will endure long after the Alliance has crumbled to dust.

You kill him.

His last words, before he dies, are his daughter's name.

And then Genn Greymane says "Rastakhan chose his fate when he sided with the Horde." And Shaw says "We can savor our victory back in Boralus."

Those lines land as hollow and ugly, and they are supposed to. Because the writing made Rastakhan sympathetic. He wasn't a villain. He was a king defending his home. And you just killed him.

Then the perspective flips. You become the Horde. You see what you did from the other side. And then you hunt Jaina.

Compare this to Garrosh: a warchief who turned his capital into a military prison, who unearthed an artifact of Old God power and used it to corrupt his own people, who betrayed the Horde's leaders and murdered anyone who challenged him. There is no ambiguity in Garrosh. He earned what's coming. You are justice.

Rastakhan earned nothing except a foreign invasion he couldn't survive.

This is the moral spectrum these two raids are navigating, and it is a spectrum that most raids never touch. You can rank Ulduar by its hard modes. You cannot rank what it feels like to realize, halfway through a boss fight, that you are the villain of this story from a certain point of view. There is no tier list category for that.

I want to say something about music, and I want to say it carefully, because there is a complication.

Players can turn the music off. In *World of Warcraft*, as in most PC games, the audio settings are fully in the player's hands. This means that any design decision that relies on music being present is built on sand. You cannot make music load-bearing. The raid has to work in silence.

The Siege of Orgrimmar works in silence. The geography, the momentum, the familiarity weaponized against your accumulated experience — none of that requires sound. The emotional architecture stands without it.

But the Battle of Dazar'alor has a secret for the players whose music is on.

The orchestral suite that plays through that raid — anchored by the piece called "Zandalar Forever" — is some that I believe is among the finest music *World of Warcraft* has ever produced. It has a raw energy, a forward momentum, a quality of something being built and then released that matches the raid's narrative arc almost perfectly. And at its center is percussion. Heavy, insistent, unmissable percussion. Drums.

This is not accidental. Drums are pre-linguistic. They bypass the analytical brain and go straight to something older — heartbeat, march, urgency. You don't decide to feel the drums. You feel them. And in a raid that has already primed you with the sensation of military advance, of being the tip of something massive moving forward, the addition of that sonic layer is almost unfair in how effective it is.

Some players experienced the Battle of Dazar'alor as a good raid. Some experienced it as something they will carry for years. The content is identical. The difference, in many cases, is whether the music was playing.

This is what I mean when I call music a force multiplier rather than a foundation. It cannot create emotional resonance from nothing. But where emotional resonance already exists — where the design has already done the work of making the player feel something — music can take that feeling and make it enormous. It is additive in the best possible sense. It elevates what already works into something that transcends the category of "game content" entirely.

The players who will never know what they missed are the ones who turned it off.

I want to return to those tier lists, because I think they reveal something important.

When Siege of Orgrimmar appears on a best-raids list, the language around it is almost always narrative. "The most cinematic raid in *WoW*'s history." "Palpable stakes." "A raid that gave you the feeling of being part of something larger." When it doesn't appear, the criticism is almost always structural – it ran for too long as the only available tier, causing burnout; some of the bosses in the middle felt like filler; the subterranean Garrosh sections lost the momentum of the city assault.

These are fair criticisms. The raid is not mechanically perfect. But here is what strikes me: the criteria being used to evaluate it are the wrong criteria for what the raid is actually doing. Asking whether Siege has elegant hard mode design is like asking whether a war film has good car chases. It's a real question. It's just not the right question.

The raids that consistently top the lists – Ulduar, Throne of Thunder, Black Temple – are masterworks of encounter design. They are studied, iterated, praised for specific mechanical innovations that influenced everything that came after them. They deserve their position.

But they don't make you feel like you're at war. They don't weaponize your history with the world against you. They don't put you in the moral position of a conqueror and then force you to see what conquest costs. They don't end with a king's last word being his daughter's name.

The things that make Siege and Dazar'alor special are things that are genuinely difficult to articulate, and because they're difficult to articulate, they get underweighted in the discourse. Emotional resonance doesn't have a tier. Moral complexity doesn't have a ranking system. The feeling of momentum – of being the tip of a spear, of an invasion with real stakes and real consequences – doesn't fit neatly into a rubric built around encounter mechanics.

So the raids get acknowledged. Sometimes warmly. And then they get ranked below the ones whose excellence is easier to explain.

Some games are great because of their gameplay. Some games are great because of what they mean. And we have good language for the first category and almost none for the second.

This is not a *WoW*-specific problem. It shows up everywhere that critics and communities try to evaluate designed experiences. The games that win "best of all time" discussions are usually the ones whose excellence is most legible – the tightest mechanics, the most elegant systems, the most innovative design decisions. The games that lodge in memory and refuse to leave are often the ones doing something that's harder to name.

What Siege of Orgrimmar and Battle of Dazar'alor are doing is building emotional architecture. They are constructing conditions under which players might feel something genuine – not the satisfaction of a well-executed pull, not the dopamine of a tier piece dropping, but something closer to what you feel watching a film that catches you off guard. Grief, maybe. Moral discomfort. The specific weight of winning a war and not feeling clean about it afterward.

They use familiarity as a weapon. They create momentum that feels like more than mechanics. They give players different moral positions depending on which faction they're standing in. They fill the air with music that hits older and deeper than analysis. And they do all of this while also being functional raids with tuned encounters and loot tables and all the infrastructure that the game requires.

That is not a small achievement. It is, in some ways, a harder achievement than making a mechanically perfect encounter — because mechanics can be tested and iterated, but emotional resonance depends on a thousand variables you can't control, including what the player brings to the table from years of accumulated experience.

The players who knew the Vale of Eternal Blossoms before it was destroyed. The players whose music was on. The players who, as Alliance, felt the guilt of Rastakhan's last breath, and then, as Horde, felt the righteous fury of pursuit. Those players had an experience that no tier list has ever properly accounted for.

They were playing the same raid as everyone else. But they were in a different story.

And I would argue that getting players into that story — making them feel something they weren't expecting to feel inside what is ultimately a repeatable piece of game content — is one of the most sophisticated things a designer can do.

We just haven't figured out how to give it the credit it deserves.

Parsing Truth: Epistemic Decision-Making in *World of Warcraft*[38]

Modern *WoW* is a strange, beautiful hybrid – a video game wrapped around a data analytics platform disguised as a boss fight. On the surface, it's fantasy combat and flashy particle effects. Underneath, it's a thousand invisible decision trees, a lattice of cause and effect, inputs and outputs, knowns and unknowns.

We like to think of ourselves as rational actors. Players. Agents. But more and more, *WoW* isn't just a game of reflexes or preparation. It's a game of epistemics – of how we know what we know, and what we do with that knowledge under pressure.

That word – epistemics – is worth unpacking. Epistemology is the formal study of knowledge itself: what counts as knowing something, how we justify a belief, how we separate truth from conviction. I use it here in a more practical sense: not just what you do, but how you know what to do, and why. That distinction turns out to matter enormously in a raiding environment.

WoW players today have no shortage of information. In fact, we're swimming in it. Combat logs record every cast, hit, dodge, cooldown, and death in exacting detail. Addons give us live readouts of damage dealt, buffs applied, healing throughput. And behind it all, sites like WarcraftLogs transform that raw stream into gorgeous, digestible output – charts, timelines, comparisons, rankings. It's beautiful. It's powerful. And it stops exactly one layer too soon.

Because here's the thing WarcraftLogs can't tell you: what to do with what it shows you. That isn't a criticism of the tool. It's a category distinction. Knowing that you parsed 87th percentile this week is information. Knowing *why* you missed 90 is knowledge. Knowing how to fix it without breaking everything else in the process – that's wisdom. Most players don't make it past the first rung. They conflate visibility with understanding, metrics with mastery. But metrics themselves are not meaning, they are the beginning of an investigation, not its conclusion.

The transformation from information to wisdom requires something no tool can automate: lived experience. Pattern recognition. Intentional reflection. It requires you – the player – to become the final layer of the system. Raw logs give you everything that happened. A structured report lets you see it in useful form. Log review with context tells you *why* something happened. And targeted coaching, strategic adjustment, the actual application of that insight – that's where wisdom lives. Each step demands more of you than the last.

This is where epistemology stops being abstract and starts being an edge.

Raid performance, viewed through this lens, sorts itself into something like a hierarchy. At the base: mechanical understanding, the ability to press your buttons correctly. One step up: instructional execution, following directions when given them. Above that: strategic autonomy, knowing the fight and your role well enough to act without being told. And at the top – the level that actually separates great players from good ones – something I'd call epistemic fluency. The ability to know the difference between what must be obeyed and what can be bent. To break rules deliberately, and correctly.

[38] This essay is adapted from one I wrote for my personal blog, found here: https://aeon.ghost.io/parsing-truth-epistemic-decision-making-in-world-of-warcraft/

Players at that level seem instinctive. They're not. They're calibrated. They've internalized the heuristics so completely that they don't wait for instructions – they anticipate. They lead. And when they improvise, they do so with surgical precision, because they understand not just the rule but the reason behind it.

WoW is relentless in how it tests this. Every fight is a decision-making crucible, moment by moment, second by second. What is the best course of action *now*? What about now? Not just in theory – in real time, under pressure, before the window closes. And the hardest part isn't knowing the right answer in a vacuum. It's recognizing it when everything is happening at once, and trusting your judgment enough to act on it.

Even the best tools can't take you there. A mirror can show you what happened. It cannot tell you what matters. Was your parse low because of bad positioning? Did one player outperform another because they had a job that required downtime, or even just mental overhead? Did you hold a cooldown to avoid overlap, knowing a simulator would punish you for it? Did you sacrifice your numbers to prevent a wipe? Those are epistemic decisions. They involve weighing not just data but intent, context, and consequence – and no overlay in the world can make that call for you.

The temptation, always, is to reach for better tools. Smarter sims. More overlays. Predictive AI. But the bottleneck has never been data. It's judgment. Wisdom under pressure. Epistemic training, not algorithmic refinement. The truly great players are separated from the merely good not by reaction speed or gear score but by their ability to ask better questions about their own play, answer those questions honestly, and adapt.

That's the game beneath the game. And it's the one that never stops.

When Games Grew Up: An Ode to *Final Fantasy Tactics*[39]

Final Fantasy Tactics is getting a remaster.

The game came out in 1998 on the original PlayStation – a tactical RPG wearing *Final Fantasy* clothes, speaking with a voice that had no business coming out of that console at that moment in history. And now, a new generation is going to meet it.

That's worth writing about.

I won't spend much time on the logistics – those details are still murky and better covered elsewhere. What I want to talk about is why this game mattered. Why it still does. And why a remaster is not just a business decision or a nostalgia play, but something that feels, at least to me, like a small act of cultural restoration.

The surface description of *FFT* is almost willfully inadequate. Tactical RPG, set in the *Final Fantasy* universe, grid-based combat, turn-based. All of that is technically accurate and says almost nothing useful. What it actually is: three-dimensional chess, played on terrain that has elevation and line of sight and geometry that changes your odds before you've made a single decision. Chasms you can't cross. Walls that arrows can't penetrate. Position matters – above and behind an enemy means they'll never see you coming; in front of and below them is a prayer. Turn order is dictated by character speed, and actions have timing, which means a mage can move forward to cast a spell only for the target to act first, rush her, and cut her down before the spell resolves. Now do that with a party of six against a similarly sized enemy force, on battlefields ranging from open plains to narrow urban alleys.

I loved every minute of it.

The game wasn't made by the *Final Fantasy* team. Its real origin story begins with Yasumi Matsuno, the designer behind *Tactics Ogre* – a different tactical RPG series known for layered politics, moral ambiguity, and brutally honest narratives. Squaresoft recognized what Matsuno was doing and brought him over, along with several of his key collaborators. The pitch was essentially: *Tactics Ogre, but in the Final Fantasy universe – and make sure to include chocobos and crystals so people know what they're playing.* And so *FFT* was born: a political drama in familiar garb, speaking with a very different voice than anything the *Final Fantasy* brand had produced before.

It would have been enough to just be good. Tactical RPGs had a niche following, and *Final Fantasy* – fresh off the runaway success of *VII* – was one of the most bankable brands in gaming. A capable entry would have sold. But it was more than that. Much more. And I don't think anyone outside the development team knew just how much more it was going to be.

Let me be honest about where I was coming in. I love games because they tell stories – stories I get to participate in. They're like books or films, they build worlds and pull you in, but unlike those, games let you take part. You're not a spectator. You're inside the story. That's always been the magic for me.

I came in through *Final Fantasy IV* and *VI*. Both told sweeping tales of war and identity, with big casts and unforgettable set pieces. They had nuance – tragic villains, hard moments – but the structure was clear:

[39] This essay is adapted from one I wrote for my personal blog, found here: https://aeon.ghost.io/when-games-grew-up-an-ode-to-final-fantasy-tactics/. For context, it was written prior to the release of the remaster.

here are your heroes, there is the villain who wants to destroy the world, go stop them. Simple. Satisfying. Morally unambiguous.

What I wasn't ready for was what *FFT* did instead.

You play as Ramza Beoulve, youngest son of a powerful noble house. Your father commands a military order. Your older brothers will inherit everything. You're a cadet, barely more than a squire, sent to put down a peasant revolt. The instructions are simple: protect the nobility, keep order, don't ask questions. Riding beside you is Delita Heiral – technically the son of a lowborn servant in your household, but to you, your best friend, your true brother. He trains beside you. Fights beside you. When tragedy strikes – his sister killed in the uprising, by someone allegedly on *your* side – he leaves.

Then the world starts to zoom out.

A war of succession erupts. Two factions claiming the throne. Families fracture. Ideals dissolve. Delita resurfaces, no longer a squire but a rising player on the national stage – aligning with lords, manipulating nobles, doing whatever it takes to seize control of his own fate. He plays the game. And he plays it well.

Zoom out again.

The Church emerges. Claiming neutrality, seeking peace, decrying the war ravaging the continent. But you can see what they're doing: pulling strings from the shadows, shaping kings, rewriting history, feeding the ambitions that keep the conflict alive. The war is a stage. The Church holds the script.

Zoom out again.

The Zodiac Stones – sacred relics, bound to saints of legend. Except they're not holy. They're demonic. They house the Lucavi, monstrous beings who prey on human desire. The Church is not just complicit. It's possessed.

Ramza learns the truth. For that, he's excommunicated. Branded a heretic. Hunted by the Church. Forsaken by his own family. He does end up facing a world-ending threat, but not as a chosen hero or a figure of legend – just as someone who knows what's coming and refuses to blink.

He stops the end of the world.

Nobody thanks him. Nobody even knows.

Delita crowns himself king. The Church buries the truth. Victors write the history. Ramza fades into obscurity – unnamed, unmourned, unremembered. The game ends with Delita alone, having just murdered the queen he married to ascend the throne, wondering aloud: "What did you get, Ramza? No one will ever know..."

He knew, and it ate him alive. Because Ramza did what was right and lost everything, while Delita did what he believed was necessary – and lost himself in the process.

I was fifteen when I played this game. Old enough to crave complexity, young enough that I hadn't yet expected to find it in a video game. I thought I knew what the medium had to offer. Instead I got political betrayal, class warfare, religious manipulation, possession by forces unseen, a protagonist who was correct about everything and punished for all of it. *FFT* didn't hold my hand. It didn't offer moral clarity. It gave me questions and trusted me to sit with them – to see for myself that good men die, that history is a lie told by those who win, that doing the right thing may cost you everything and still leave the world unchanged.

It was the first work of any medium – not just any game – that respected me enough to not sugarcoat what it had to say.

Was it flawless? No. Brilliance rarely is. The original translation was clunky and sometimes incomprehensible. The difficulty curve was steep and uneven. Recklessness – or bad luck – could permanently cost you characters, or hours of progress if you hadn't saved often enough. The polygonal graphics and 2D sprites have not aged gracefully. But the musical score remains one of the best in the history of the medium, and I still listen to it. The job system was endlessly inventive – you could build your party your way, and no two players built the same team. The world of Ivalice was rich and layered and lived-in. There have been re-releases, and a sequel aimed at a younger audience that softened the edges in ways that weren't for me. But that's fine.

Because now we're getting a real remaster. A chance to bring this story forward with the polish and reverence it deserves.

FFT showed me that video games could be art. That they could be political. That they could be honest – willing to say something difficult and trust the audience to receive it. It didn't pretend neutrality. It had something to say, and the courage to say it.

That's why the remaster matters. Not for nostalgia. Not just because it's a classic. But because this game still has something to teach – about power and truth, about what it costs to do the right thing when no one is watching, when no one will ever know.

Ramza knew.

That was enough.

What's Left When Meaning Breaks: A Look at *Nier Automata*[40]

During the raid beta test for *Mists of Pandaria Classic*, I found myself standing in the Heart of Fear, watching Imperial Vizier Zor'lok fill the room with swirling patterns of orbs. On the hardest difficulty, multiple overlapping spirals traveled at different speeds, and the personal mandate was simple: dodge, or die. I was struck by a nagging sense that I'd seen this before – and not because I remembered the fight from *Mists'* original launch in 2012.

Oh, I thought. This is a bullet hell fight.

Bullet hell – or *danmaku*, the Japanese term, meaning roughly "bullet curtain" – is a subgenre of shoot-'em-ups where the player faces such overwhelming volumes of projectiles that survival becomes an exercise in pure spatial reading. Our first significant crossover from Japan was Ikaruga in 2001, critically admired and commercially catastrophic: under seven thousand copies sold in its first week in the US, barely thirty-three thousand total across all platforms including Japan. The numbers were bad enough that the genre seemed to have answered its own question about Western appetite. The only way another bullet hell game was going to get made was if a developer decided the market's preferences were, essentially, someone else's problem.

Enter Yoko Taro.

Taro is the director of *Nier: Automata*, and he is – there is no more precise way to say this – an odd duck. He wears a creepy moon-headed mask in interviews. He insists he's just there for the paycheck. He is probably lying. The original *Nier*, released in 2010, was the kind of game that critics didn't quite know what to do with and almost no one bought – a weird, genre-bending, emotionally devastating mess with a stunning soundtrack, wildly ambitious storytelling, and combat that was, charitably, kind of janky. A beautiful failure. Beloved by the few who played it.

A sequel seemed unlikely. It happened anyway. This time, Taro handed the combat to PlatinumGames – the studio behind *Bayonetta*, slick stylish action as a matter of house philosophy – while he retained the writing, the directing, and the metaphysical trauma. What came out the other side wasn't an improvement so much as a reinvention. Nier: Automata released in 2017 to critical and commercial success. It sold millions. It won awards. It launched anime, novels, crossovers. It was unexpected, and it was unqualified.

The premise sounds straightforward enough. Thousands of years from now, humans have fled to the moon, Earth overrun by alien-built machines. The only ones still fighting are YoRHa androids – purpose-built humanoids, sleek and tactical. You play as 2B, a combat android who looks like a fetish model (blindfolded, thigh-high leather boots, black miniskirt with a slit up to the waist, a katana that is frankly unreasonable in scale), accompanied by 9S, a scanner android who looks like a prepubescent schoolboy (also blindfolded, shorts, an odd double-breasted coat, a much smaller katana). Your mission: exterminate the machines, win the war for your human masters.

You start doing exactly that.

---

[40] This essay is adapted from one I wrote for my personal blog, found here: https://aeon.ghost.io/whats-left-when-meaning-breaks/

And then the anomalies begin. Some machines are peaceful. Some have formed small villages, primitive societies. A few have begun worshipping gods of their own invention. They're the minority – most machines are still hostile – but the more time you spend among them, the more the line between android and machine begins to blur. The question the game plants early and never stops pressing: who is the real enemy? And underneath that, the sense that you've been here before – that you're moving through a story that has happened, that is happening, that will happen again.

That sensation is intentional. *N:A* structures itself as Routes. Finish Route A and you might assume you're done. You're not. Route B shifts your perspective, fills in context you didn't know was missing. Route C collapses the structure entirely. Routes D and E reframe your role as a player in ways I won't describe here, because the discovery is the point. The same story beats return each time, but viewed from a new angle, with new knowledge – what you thought was one thing reveals itself to be something else, and then something else again. Most games' plots are a fixed sequence of events, like a script: you follow them to the end, credits roll, done. Play *N:A* that way and you'll see less than a quarter of it.

The music works the same way. It was the City Ruins theme that made me understand what the soundtrack was actually doing. Walking through the ruins, you hear the standard version – haunting, melodic, immediately distinctive. Step inside a building and it becomes muted, echoed, intimate. Enter combat and it gains urgency, the tempo alive with threat. Hack into an enemy and it fragments into glitchy chiptune. Through every shift, the melody stays the same, and the playback position stays the same, so the transitions feel not like track changes but like recontextualizations – the game running parallel versions of every song simultaneously, ready to swap instantly without breaking emotional continuity. The music isn't atmosphere. It's architecture. It is the first soundtrack that made me rethink what a video game soundtrack is capable of doing.

The philosophy of the game is harder to summarize without spoiling it, and I'm deliberately not going to try. What I'll say is this: 2B's stoicism isn't coldness, it's armor, forged by a repetition the game reveals slowly. 9S is driven by a desire for truth and understanding, and finds that the more he learns, the less anything means. The questions accumulate in sequence. First: what does it all mean? Then: how do I know what's real? Finally: what's the point of it all? The game blends existentialism with nihilism and sets the whole thing in a world that doesn't reward goodness – not because it's cruel, but because it's indifferent. Which somehow hurts more. It's not philosophical because it name-drops Kierkegaard, or because the main character's name is literally a reference to a core existential question. It's philosophical because it makes you feel those ideas. There's a meaningful difference.

I don't have a clean ending here. That's fine. *N:A* didn't have one either.

The game's finale deploys bullet hell as metaphor made literal. As the credits roll – set to one of the best songs in the entire soundtrack – you play as the small spaceship you've been piloting during hacking minigames throughout the game. The bullet curtain grows progressively denser, more overwhelming. By the end, it doesn't matter how skillful you are. You cannot win. The screen fills. You die.

The game asks if you want to give up. You say no.

You try again. You fail. You try again. You fail. You try again. You fail. And at the lowest point, something changes. You hear a voice. Then another. The music opens into a chorus. And the little cursor

spaceships of other players – the ghost data of everyone who made the same choice you're making now – begin to fly beside you. They intercept bullets. They shield you. They save you.

You make it through.

The credits end. The game offers one last choice: are you willing to delete your save data to help a stranger complete this moment? You will lose everything. But someone, somewhere, will be saved.

You are prompted to write a small message. You do. And then your story ends, so that someone you will never know, someone you will never meet, can have what you just had.

It has been eight years since I played it. I finished all the Routes and haven't touched it since, and I think I know why: I'm afraid of diluting it. Of returning to something that hit with that kind of force and finding it smaller than I remembered. Some things you only get to discover once.

What it gave me was questions. Memories. The haunting sense that meaning isn't found – it's built, and sometimes it breaks, and what you do in that breaking is the only answer the game is willing to offer.

This game isn't about winning.

It's about who you are when winning is no longer possible.

Dramatis Personae

Aeon Aliases

Battle.net — *StarCraft / Brood War, Warcraft III, Diablo II*

- kommpakk
- winter.mute
- Wintermute[cF]
- Winter[cF]
- ce[n]arius

Battle.net 2.0 — *StarCraft II, Diablo III, Heroes of the Storm*

- Wintermute
- Winter
- Aeon#12420 *(current — friend requests welcome)*

World of Warcraft — Original Timeline (all levels noted as pre-squish)

Sen'jin

- Enma — Orc Shaman, Horde, level 60
- Amne — Human Mage, Alliance, level 30

Korgath

- Retcon — Blood Elf Paladin, Horde, level 70
- Chairmanmeow — Tauren Druid, Horde, level 70
- Burninate — Blood Elf Warlock, Horde, level 70
- Miluda — Blood Elf Rogue, Horde, levels 70–80

Khadgar

- Aeonwyn — Night Elf Rogue, Alliance, levels 80–110
- Aeonwynn — Human Paladin, Alliance, levels 80–110
- Aeonith — Night Elf Hunter, Alliance, levels 80–110
- Aeonista — Human Priest, Alliance, levels 85–110
- Aeonix — Draenei Shaman, Alliance, levels 90–110
- Aeonlysa — Human Mage, Alliance, levels 90–110
- Aeonfyre — Human Warlock, Alliance, levels 90–110
- Aeonyx — Human Death Knight, Alliance, levels 100–110
- Aeonna — Night Elf Druid, Alliance, levels 90–110
- Aeonari — Night Elf Demon Hunter, Alliance, levels 100–110

World of Warcraft — Classic Timeline

Pagle

- Aeonwyn — Human Rogue, Alliance, levels 60–90
    - *Individual Achievements*
        - Temple of Ahn'Qiraj — World 90th (Combat)
        - Black Temple — World 19th (Combat)
        - Dragon Soul — World 44th (Assassination)
- Aeonwynn — Human Paladin, Alliance, levels 60–90
    - *Individual Achievements*
        - Icecrown Citadel — World 163rd (Protection)
        - Mogu'shan Vaults — World 73rd (Retribution)
        - Heart of Fear + Terrace of Endless Spring — World 40th (Retribution)
        - Throne of Thunder — World 9th (Retribution)
- Aeonith — Night Elf → Worgen Hunter, Alliance, levels 70–90
    - *Individual Achievements*
        - Firelands — World 46th (Survival)
        - Dragon Soul — World 178th (Survival)
- Aeonna — Night Elf Druid, Alliance, levels 60–90
- Aeonista — Human Priest, Alliance, levels 60–90
- Aeonwin — Human Rogue, Alliance, levels 60–90
- Aeonstorm — Draenei Shaman, Alliance, levels 70–90
- Aeonir — Human Mage, Alliance, levels 70–90
- Aeonyxi — Human Warlock, Alliance, levels 70–90
- Aeonwinn — Draenei Paladin, Alliance, levels 70–90
- Aeonxi — Human Death Knight, Alliance, levels 80–90
- Aeontiny — Gnome Rogue, Alliance, levels 80–90
- Aeonari — Human Warrior, Alliance, levels 80–90
- Aeonyth — Night Elf → Worgen Hunter, Alliance, levels 85–90
- Aeoncutie — Worgen Hunter, Alliance, levels 85–90
- Aeoness — Worgen Hunter, Alliance, levels 85–90
- Aeonesque — Worgen Death Knight, Alliance, levels 85–90
- Aeoniel — Draenei Paladin, Alliance, level 90
- Aeonael — Human Paladin, Alliance, level 90
- Aeonhua — Pandaren Monk, Alliance, level 90

Acknowledgements

While I consider *The Aeonomicon* to be a work of a single person, no man is an island, and I am no different. There are many people without whom I would not be the person I am today, or be able to produce what I've done here. They deserve my profound gratitude, and I will do my best to live up to that.

To Clover: I am eternally grateful for your love and support, and for teaching me – over and over, and with patience I don't always deserve – what those words actually mean. Thank you for being my person, for seeing me clearly enough to remind me of who I am, and for making that person someone I'm proud to be.

To the leaders who gave me the privilege of just being a player: Chains, Mudflaps, Stur, Kast, Aradell, Jordan, Panda, Funkey, and Dot. You carried the burden of leadership so I didn't have to, and you did it well enough that I rarely had to think about it. I learned much from watching each of you lead – directly and by example – and each of you had distinctive lessons to teach. Thank you for your patience. Thank you for the lessons.

To everyone who played for Starcaller – *Cerein, Cloud, Crazi, Farshin, Frick, Gloriosum, Iburysox, Konstance, Lothaestus, Martharion, Phiesta, Stormy, Tastykake*, Bazze, Bobtheknight, Borkborkbork, Chiaku, Crunch, Drezni, Eien, Emiwoo, Epsilon, Escoo, Hajek, Kai, Matikz, Moranall, Nuleen, Okearney, Pearalta, Pewpewkachoo, Pippen, Pompompoms, Quorum, Rogerrabblt, Rollo, Selrenor, Sinsational, Sporkfed, Yttrium, and Zarol: you agreed to participate in my mad science experiment, you followed where I led, and you gave everything you had while I was still figuring out what I was doing. What we built together – and what we earned – I carry with genuine pride. Thank you for trusting me with it. Thank you for giving me the best of yourselves.

To Chris Metzen, Samwise Didier, Rob Pardo, Jeff Kaplan, Mike Morhaime, Allen Adham, Dustin Browder, Glen Stafford, Matt Uelmen, Glenn Rane, Peter Lee, and Justin Thavirat: there are many names in the Blizzard pantheon, but yours are the ones that come to mind first. You don't know who I am. I know exactly who you are, and what you made has been a critical piece of making me who I am. Thank you for having a vision, and for turning it into something real.

To David Howe: thank you for inspiring this project. May your journeys be joyful and I wish you great success in your new endeavours across the Atlantic.

To everyone at Blizzard – past and present – who has been part of keeping these stories alive: thank you.

About the Author

Donald Coles is a program manager and delivery leader whose qualifications, professional history, and gaming opinions are documented at considerable length in the preceding pages. He identifies as a paladin, though he is willing to dabble at being a rogue or a hunter as needed. He currently lives in Houston, Texas with his wife, daughter, and three excitable parrots.

About the Book

The Aeonomicon was written over the course of two months in the spring of 2026, drawn almost entirely from personal memory.

Several essays in the end matter existed in prior forms and were revised and integrated for this edition; the remainder were composed from scratch. There were no research trips and no archive of notes to consult – only recollection, considered at length and committed to the page.

Every stage of production – conception, drafting, revision, copy editing, layout, printing, and final assembly – was completed by the author alone. What you are holding is, in that sense, a genuinely handmade object: a single person's sustained effort to render a life's work into something you could hold. This is the First Edition. Each copy of this printing is unique.

The epigraph is from the poem "Ulysses," by Alfred, Lord Tennyson.

If this book has reached the end of its journey without finding a home, the author would be glad to have it returned. He can be found at dlcoles@gmail.com to make arrangements.