

Reflecting on 30 Years of Global Technology Leadership

A person in a dark suit stands with their back to the camera, looking out over a vast, glowing circuit board that stretches towards a city skyline in the distance. The scene is set at night, with the city lights reflecting on the surface of the circuit board. The overall tone is futuristic and technological.

A journey from Netscape Navigator to the age of
autonomous AI agents

"I was a translator of technology to business."

By **George Ralph**

[LinkedIn: George Ralph](#)

[X: @georgenwralph](#)

INTRODUCTION

"Well, this is different."

After 30-plus years of helping businesses grow using optimal technology solutions, I'm taking a well-earned sabbatical. A time to temporarily hit pause and reflect before I throw myself into the next major challenge.


I have a natural instinct to want to fix things; to understand how things work, including people. The signs were there at a young age. When I was eight, my parents' washing machine broke. Before they could do anything, I'd already taken it apart to uncover the problem and fix it.

Guess what... *I studied psychology at college.*

For now, I've enrolled to do the UK Government's NED course to support small businesses up to FTSE 100. I'm also preparing to complete an AI practitioner course to become a certified government advisor on AI usage and I am acting as a volunteer.

But for the rest of this piece, I want to share my reflections on what I've seen and learnt as a technology / business leader. Having grown up before mobile phones and the internet, and experiencing the explosion of technology innovation throughout my professional career, it has been a unique time to witness Moore's Law in action. The world has changed at a breathtaking pace.

30+ years · 4 MSPs built · From 14kbps
dial-up to AGI



It All Began in 1994

I've spent most of my career understanding businesses, processes, the bottlenecks created by people and then using technology to help create a competitive advantage. In short, I brought technology to people for whom access was an issue.

I was a translator of technology to business.

Nurturing this skill set began in 1994 when I started my career teaching people to use Netscape Navigator, the popular web browser which launched that December. This was the year when dial-up modems doubled in speed from 14 to 28 kbps, Quentin Tarantino released Pulp Fiction, Friends premiered for the first time, and the Sony PlayStation launched.

Those pre-social media, pre-algorithm days taught me an important early lesson: technology is only valuable when people can understand it and use it confidently.

What seemed intuitive to someone working in technology could feel completely alien to a user. I learned to listen, remove jargon and explain the purpose of the technology rather than simply describe how it worked.

This particular focus for me really helped during my CTO and Tech Director days shifting firms from OS2 to NT4 and the emergence of graphical interfaces (which all IT folk thought would be the end of our jobs) – this was the pivot in helping give access to technology for everyone, not just us text screen geeks!

That stayed with me as I moved into systems administration and later into leadership.

Whether you are introducing a new platform, developing a cybersecurity strategy or discussing AI with a board of directors, the challenge fundamentally remains the same: you have to connect the technology to a real human or business outcome.

During the 90's, as I moved on from web browser training, Infrastructure projects, and other engaging tech mega changes including a government body, I went on to form or join what would become the first of four managed service providers (MSPs).

And soon my role evolved.

Increasingly, my focus was on connecting technology to commercial outcomes, risk management and organisational confidence. Back in those days, and even moving in to the noughties, most businesses I dealt with had a handful of IT staff; sometimes just one or two. As a technology leader and CTO, I was the last line of defence.

I could never switch off.

I loved it.

Before the explosion of AI, which has drastically widened access to technology – we all carry a mobile computer in our pockets – it was down to me to stay ahead of the technology curve, learn what the latest developments were, and apply them appropriately for every client.

That is still true today but the focus has shifted. It is not so much translating technology to meet business needs, and persuading CEOs of its virtues, but persuading technologists that the business needs them. After all, there are so many technology jobs. There's always an attractive alternative; better pay, better equity options, better product innovation etc. Talent retention is highly pressurised for many businesses.

- 1 — 1994
Teaching Netscape Navigator. Dial-up modems hit 28kbps.
- 2 — Late 90s
Systems administration, infrastructure projects, government bodies.
- 3 — Early 2000s
Formed or joined the first of four MSPs.
- 4 — 2000s–2010s
Outsourced CTO/CISO/COO across global clients.

Technology is only valuable when people can understand it and use it confidently.



Clarity Helps Frame the End Destination

Building businesses teaches you that there is nowhere to hide. Simultaneously, you have to understand: the technology, the client, the finances, the people/internal culture, and the commercial model.

When I look back on the early years, I would use open source technology (CentOS, Redhat, Ubuntu) Linux systems, and combine them with things like VMware to build virtual machines and compile the software to create virtual desktops. It kept us honest. There was no 'off the shelf' cloud environment to leverage.

One of the most important lessons I learned was that good technology alone does not create a successful company. You need a clear proposition, disciplined delivery, strong client relationships and a team that understands what the business is trying to achieve.

Being an outsourced COO, CISO and CTO, I have had the privilege of understanding how a myriad of different organisations work from the inside. Over three decades, this has allowed me to develop my own internal pattern recognition, identifying business constraints resulting from clients' operating models, suppliers, or technical decisions.

One of the key lessons I've learned from a leadership perspective is that you have to create clarity. People perform at their best when they understand the destination, the reason for going there and the role they play in getting there.

Clarity comes from listening to the client before selling solutions. My ambition has always been to act as an extension of each client's organisation and a trusted adviser to its leadership. This has served me well. It's helped me to build strong local teams that understand the importance of service standards, clear communication and accountability.

People perform at their best when they understand the destination, the reason for going there and the role they play in getting there.



Don't React. Anticipate.

One aspect of being an outsourced CTO that I've loved is the kinetic nature of technology management and finding elegant solutions to complex problems. The pace of innovation over the last 20 years has been staggering. Traditional MSPs were expected to keep systems running, resolve tickets and maintain infrastructure. That is no longer enough.

A modern provider needs to understand an organisation's investment strategy, growth plans, regulatory footprint and risk appetite. It must be able to translate those factors into an appropriate technology and security operating model. In addition, organisations expect their provider to help them interpret regulation, manage cyber risk, oversee third parties, develop data strategies, support due diligence and contribute to business planning.

The best providers do not simply respond to problems. They help clients anticipate what is coming next and make better decisions before those problems arise. In other words, they aim to be less like a traditional helpdesk and more like the 'precogs' in Minority Report; identifying potential problems before they become incidents. When I worked as an outsourced COO to MSPs this was a common feature of change I would implement.

Ultimately, when technology providers focus only on systems and tickets, they risk becoming disconnected from what the client is actually trying to achieve.



LESSON 3

To Scale, Don't Rely on Absolute Control

In my experience, you cannot scale a company purely through process. Process creates consistency, but culture determines how people behave when the process does not provide an obvious answer. Throughout my time working at, and scaling various MSPs, it became increasingly clear that in order to establish and maintain shared standards, you had to have the right balance of local market expertise, local leadership, alongside a unified commitment to think and act globally.

Every client is different. No matter where in the world they are located, they expect service quality consistency that reflects regional expectations and adheres to regulatory requirements.

Strong leadership relies on not retaining control over everything. Build the team, give people individual autonomy and genuine responsibility. And create a supportive environment where they can make good and bad decisions.

To scale effectively also requires altering the cadence of communication. For smaller companies in the early stage of growth, information flows more naturally because teams are smaller. Everything's more immediate.

As companies scale, and teams grow bigger and more complex, technology leaders have to be more deliberate, increasing the frequency of explaining the strategy and the rationale behind business decisions.

Enter AI, Scaling support and Automation.

Culture determines how people behave when the process does not provide an obvious answer.

AI: A Double-Edged Sword

I'm truly excited about the next chapter of my career. Having personally grown as a leader alongside the MSPs I've worked for and the plethora of global clients I've supported, through major technology turning points – the roll out of mobile phones, cloud architecture, cybersecurity management, Big Data, quantum computing, generative AI – I'm expecting the years ahead to be equally fascinating.

Like the kid that dismantled his parents' washing machine, the scope of what might be possible, as a technologist, is limited only by imagination.

In the AI landscape, we've recently seen Anthropic's creation of Project Glasswing to allow critical infrastructure providers and big tech to safely use Claude Mythos Preview to identify and remedy security vulnerabilities and, hopefully, stay one step ahead of cyber criminals.

The power of Mythos was clearly enough to convince Anthropic's leadership that this needed to be handled with care.

And this is the issue with AI. Like fire, it can be used for constructive (cooking food) or destructive purposes (destroying property).

One of my many concerns, is people are using tools like Claude Code to develop software but aren't fully considering who owns the intellectual property. Might we see hedge funds who use vibe coding (a term coined by Andrej Karpathy in 2025) to develop trading algorithms getting switched off because they don't own the IP? Or experiencing the fallout of a catastrophic cyber breach directed at Anthropic or OpenAI, as the AI arms race continues? Also, if you're a hedge fund CTO, how do you monitor the code for accuracy and ensure there is no hidden malware contained deep within? I know how I do it, but I'm concerned that others do not.

There are plenty of unintended risks that are yet to fully play out.

If you're a hedge fund CTO, how do you monitor the code for accuracy and ensure there is no hidden malware contained deep within?

The Rise of Autonomous AI Agents

If I were to nail my colours to the mast, the one thing technology leaders are most underestimating is the emergence of autonomous AI agents.

The conversation is often focused on generative AI producing text or assisting employees with certain procedural tasks. The more significant development will be agents that can access systems, analyse information, make decisions and initiate actions across multiple systems.

Many firms are treating AI primarily as a data-protection or employee-policy issue.

Those things matter, but the risk is becoming much broader – like a spider web effect.

An AI agent is not simply another piece of software. In some circumstances, it begins to resemble a digital member of the workforce, with an identity, permissions and responsibilities.

That raises questions about ownership, accountability and access control.

Firms need to know which agents exist, what they can access, who authorised them and how their activity can be reviewed. I would argue the first step towards preparing for this inevitable rise of orchestrated agents is for organisations to extend their identity governance, privileged-access management, auditability and segregation-of-duties principles to AI agents.

Waiting until after agents are widely deployed will make the governance challenge considerably harder.



Inventory all AI agents

Know which agents exist and what they can access

Named ownership

Establish clear accountability for each agent

Risk classification

Classify use cases according to risk level

Access controls

Apply appropriate intervention thresholds and human approval gates



SHADOW AI

Watch the Shadow (AI)

Shadow AI is another significant concern. Employees will use tools that help them work faster, particularly when the approved alternatives are too slow or restrictive. Simply banning those tools will not solve the problem.

The practical starting point is visibility.

Firms should create an inventory of AI tools and agents, establish named ownership, classify use cases according to risk and apply appropriate access controls.

Higher-risk actions should have clear intervention thresholds or human approval. AI governance should support responsible adoption, not prevent it. The objective is to create enough control and accountability for the organisation to innovate confidently.

AI governance should support responsible adoption, not prevent it.

An AI-Enabled CISO

If I were to start an AI company today, I'd design an interactive IT security companion that shadows and learns what you're doing and gives you tips on how to stay secure... 'I see you've just downloaded a productivity app, which contains a privacy notice. Can I confirm you are happy with your data residing here?'

Basically a virtual George Ralph without my WhatsApp exploding daily.

Or more prosaically, an AI-enabled CISO in a business context.

This would help address the shadow AI privacy-related risks that, unfortunately, organisations – and we as individuals – will increasingly need to contend with.

This is no ordinary paradigm shift that we are living through.

The push towards artificial general intelligence (AGI) is creating an age of hyper-connectivity. Look at the latest news that we only have 10 years left to live before AI takes over!


Soon, AI will be writing its own AI, securing AI, checking AI, agents will increasingly design newer version of themselves. What that means for technology adoption and optimisation in the financial sector over the next decade is the \$64K question.

But one thing is absolute: technology leaders will need to use it responsibly.

Those that do will become beacons of light.

This is no ordinary paradigm shift that we are living through.





WHAT'S NEXT

Bring It On.

As I look ahead to the next chapter of my career, I'm reminded once again of that broken washing machine in pieces on the kitchen floor, ready to be re-assembled.

**There will always be
another challenge.**