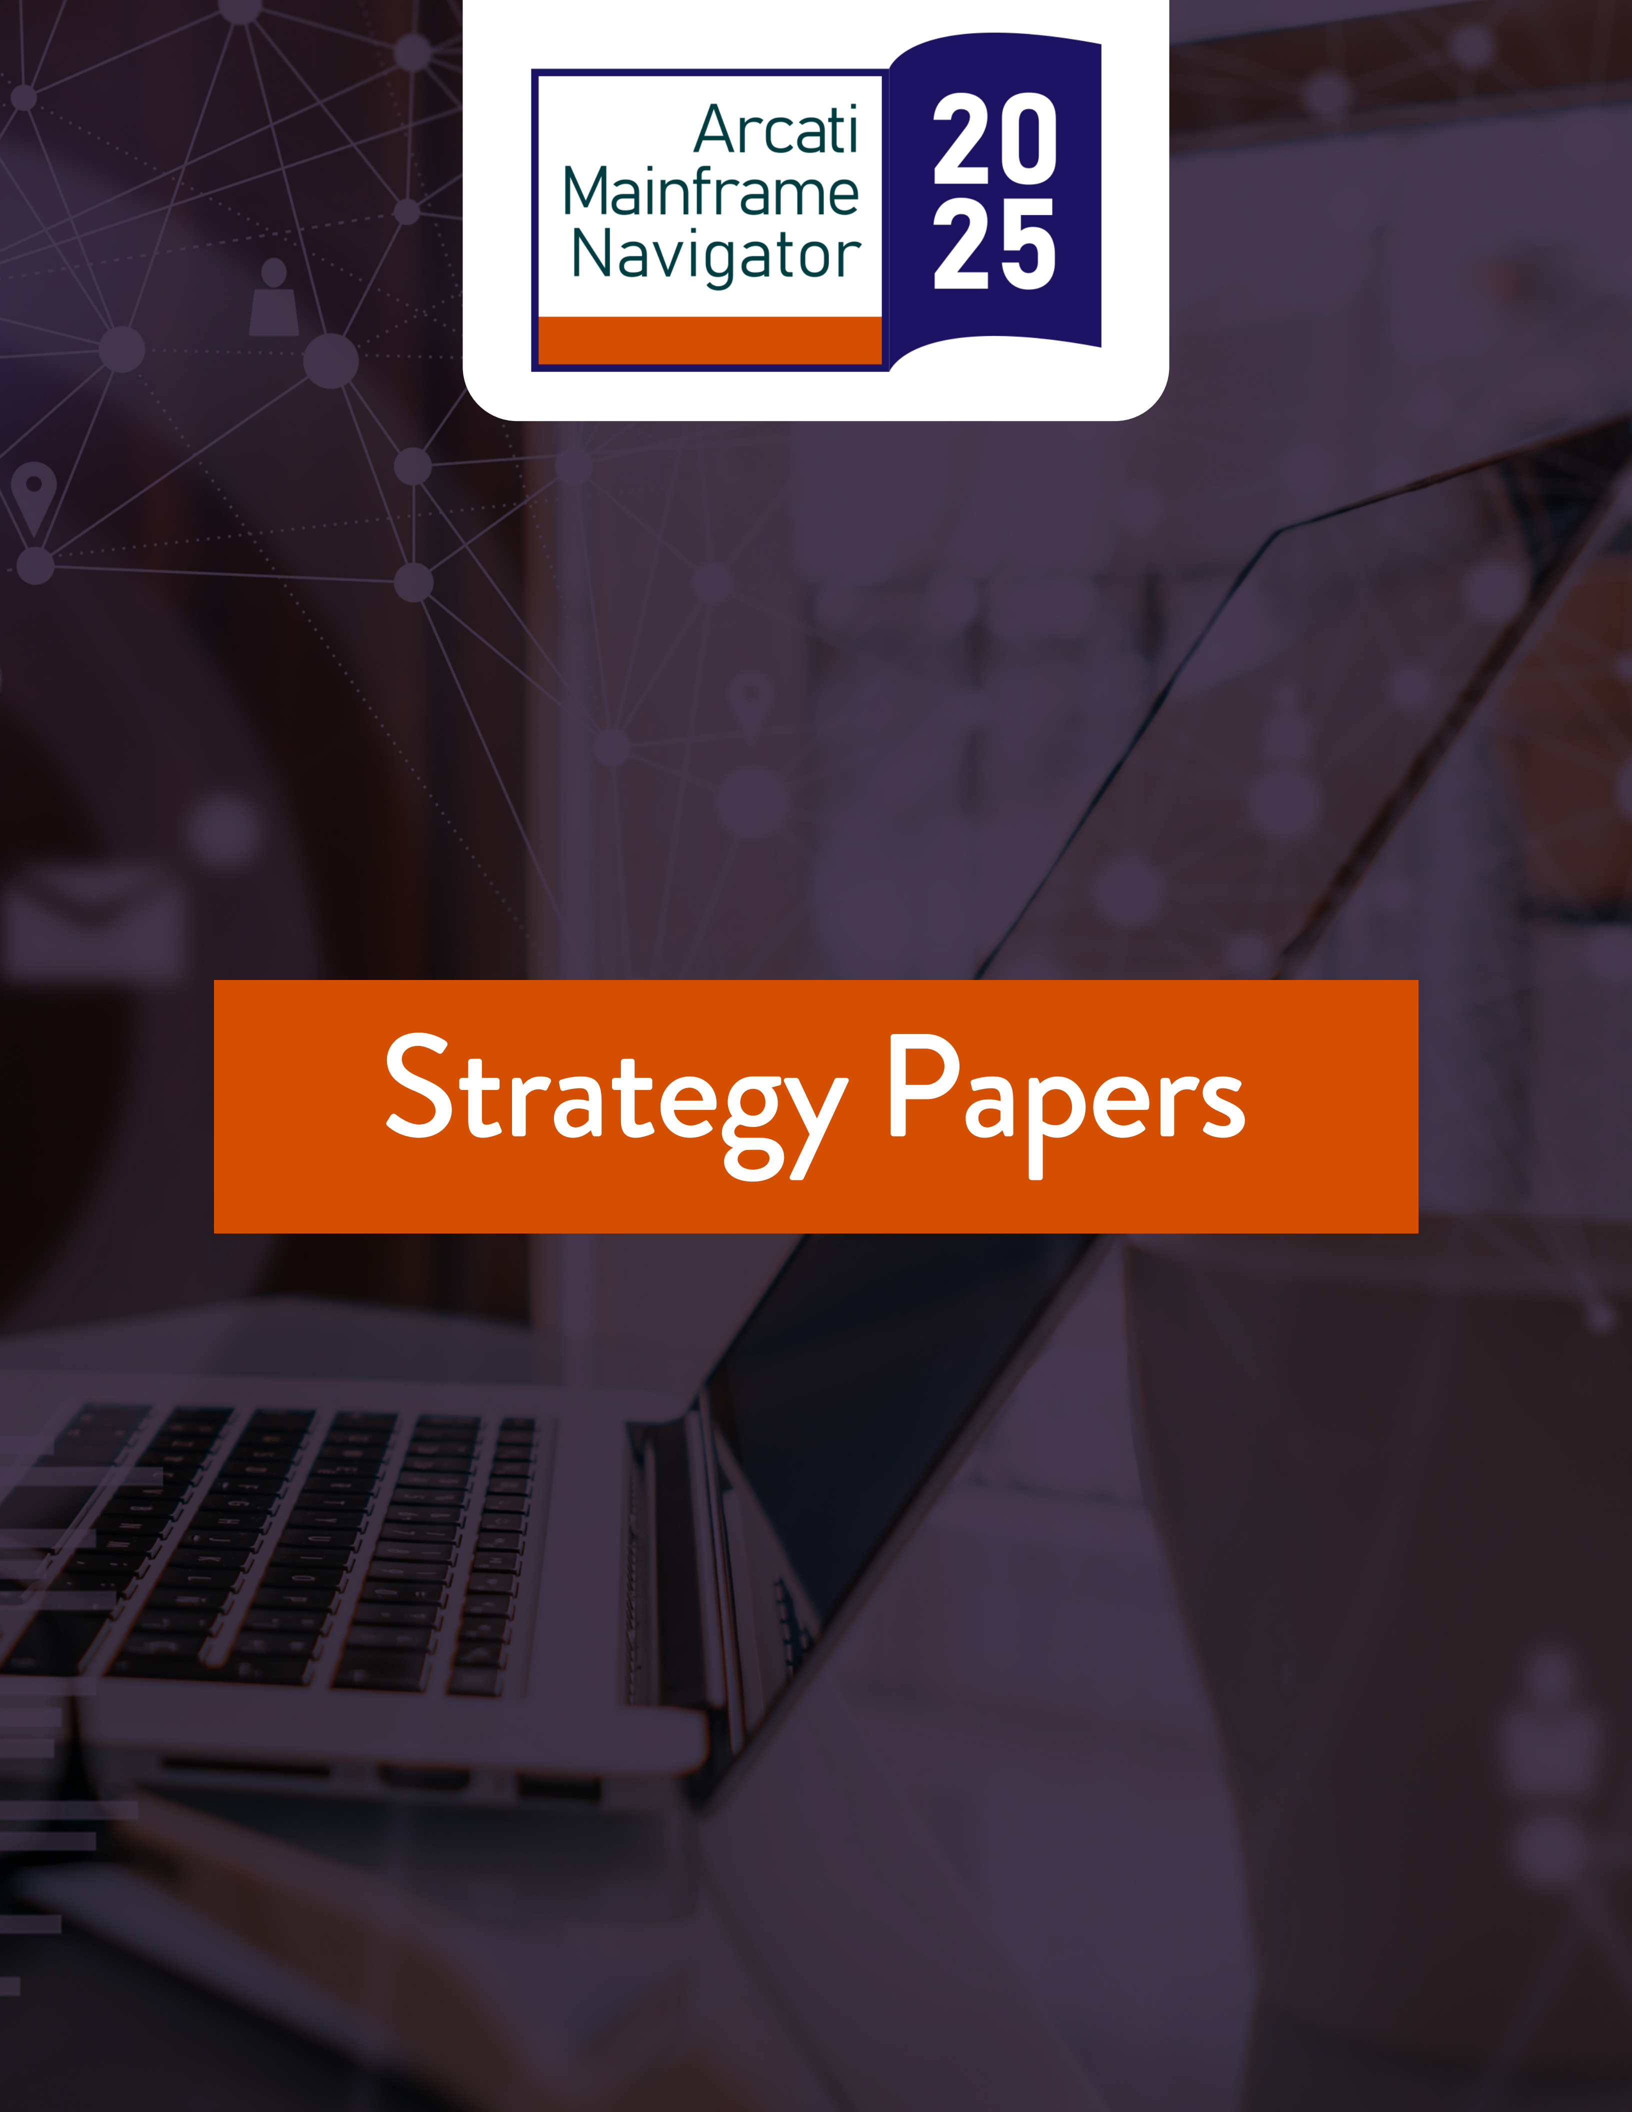

Arcati Mainframe Navigator 2025

Strategy Papers

Mainframe Security in 2025:
Countering New Threats, Using AI, and Getting the Basics Right

Introduction

As we moved into 2025, two trends were front of mind for many of our clients: their continued efforts to achieve a Zero Trust security stance in the face of an evolving threat landscape, and the unstoppable rise of artificial intelligence – countering the risks of AI while seizing the opportunities it presents.

Both trends are linked to the fact that, for many organizations, the mainframe has been seriously overlooked in the past as a cybersecurity risk. And yet vulnerabilities clearly exist, from flaws in code and configuration that can be exploited by criminals, through supply chain and vendor product risks that expose the organization, to insider threats. In short, the threat is real, and growing. IBM's Cost of a Data Breach Report 2024 found that the global average cost of a data breach in 2024 was USD 4.88 million, a ten percent increase on the previous year, and the highest total ever.

Slaying the cyber beast

The cyber threat landscape is continuing to evolve, powered by creative criminals, new tech, and a connected world. It might be likened to the mythical Greek monster, the Hydra: cut off one head and two more cyber threats spring up in its place, presenting new and more complex challenges. At the same time, the mainframe is now mainstream, as much a part of enterprise IT as anything else. And the platform is, of course, hackable. Vertali's mainframe security team has done so many times, as part of formal security assessments and penetration testing, all be it in a safe and secure manner. But once a bad actor gains access, the resulting damage can be catastrophic. Creating backdoors means attacks can escalate even after the initial threat is eliminated.

AI and quantum computing - the latter for complex concurrent computations and so cracking digital encryption faster - add to today's threat complexity, bringing the potential to assist security professionals as well as new opportunities for hackers. Recent technologies bring new back doors and vulnerabilities; today's mobile apps can significantly compromise cybersecurity, with smartphones a target for malware, and therefore another potential route into the mainframe (alongside, in a connected IoT world, things like fridges and exercise bikes, which have even less in the way of security controls). Increasingly, mobile platforms are used to access and process sensitive data: personal and corporate email, messaging services, corporate and financial data, and more.

This all requires robust policies, governance and technology-driven oversight, for both bring-your-own devices and corporate-provided tech: clear policies plus proven mitigating technologies to protect against breaches and data loss, and including logging and monitoring of all devices. And this is just one tiny piece of a hugely complicated jigsaw.

The double-edged sword of AI

It's everywhere and it's growing, across all areas of business and increasingly the public sector. In January 2025, for example, the UK government launched the AI Opportunities Action Plan to improve efficiency and have an impact on areas ranging from healthcare and education to improving roads and supporting small businesses.

Generative AI brings new risks and challenges as well as opportunities. In our world, on the plus side, we're already seeing rapid developments in areas such as Enhanced Threat Detection and Prevention, Automated Security Responses, Enhanced Encryption and Data Privacy, and Vulnerability Management.

On the debit side, AI is already used by criminals to identify vulnerabilities, develop more sophisticated phishing attacks, and automate the exploitation of security flaws. Mainframe security strategies have to evolve fast to counter these threats, but integrating AI technologies and tools into mainframe environments can be a complex ask, and can require significant investment in time, people and resources. We must also be wary that over-reliance on AI and automation does not lead to complacency, with human oversight overlooked.

So, what can we expect from AI?

In enhanced threat detection and prevention, we can look to AI-powered security analytics: analyzing vast amounts of data, identifying patterns, and detecting anomalies that may show a security breach. That can lead to more proactive and adaptive security measures, Generative AI can also automate intelligent security incident responses, and reduce the time lag between detection and targeted action - isolating affected systems and automatically applying patches and updates. AI can be used to dynamically adjust access controls based on real-time assessments of user risk, ensuring only authorized and verified users can access critical mainframe resources. In vulnerability management and predictive maintenance, generative AI can be used to analyze system behaviors, code and configurations to anticipate and predict vulnerabilities before they can be exploited. And so on.

The changing nature of AI-driven security

As AI becomes more commonplace in mainframe security and operations, we will need to ensure that our usage complies with the necessary regulatory standards and requirements. We will have to ensure transparency and auditability for security-related actions taken with AI involvement. And in terms of the new skills needed, we will see a rising demand for professionals with expertise in both mainframe environments and AI, and ideally security. Continuous training and ongoing professional development will be crucial for our security teams to keep up with, and benefit from, evolving AI technologies and approaches.

Remember the basics: mainframe security 101

In this short paper, we've taken a look at the threat landscape, and at the risks and opportunities presented by AI. It's worth adding this third element to any discussion of strategy. With continuous innovation and new functionality, such as AI and extending to developments such as Pervasive Encryption (PE), multi-factor authentication (MFA) and file integrity monitoring (FIM), it's important that we don't forget the basics, and suffer from what a colleague calls "shiny object syndrome". Buildings are only as strong as the foundations they are built upon, and it's the same in mainframe security. We need be on top of the detail and have those basics in place if we're going to achieve a true Zero Trust stance. We have a blog on this topic, but I can summarize the main points here.

Authentication - logging on or connecting to the mainframe. Accountability can only be assured if you can be sure that whoever or whatever is accessing your system is who or what they say they are. Are your password rules strong enough? How often do you require a password to be changed. Have you implemented the RACF encryption algorithm KDFAES? If you use MFA, how is it used?

Access management - the insider threat allied with phishing and ransomware pose a serious risk; no amount of encryption will prevent someone with valid access to your data from editing, copying, or selling it to the highest bidder. They may even re-encrypt it using their own key. Issues to consider in access management include:

- **Data ownership** – do you know who owns what data on your system?
- **Role-based access control** – role-based security to restrict access to authorized users, implementing mandatory or discretionary access control, is important.
- **Approval process** – all requests should be subject to some level of approval before being actioned, with different levels of access, or access to sensitive functions/data, requiring different approvals.
- **Automated systems** - can complicate issues, requiring other factors to be considered, including issues around Privileged Access, and the principle of Least Privilege.
- **JML** – redundant accounts and access can provide a way to bypass security controls. Rigorous JML processes should reduce the risk, provided they are followed.
- **Recertification** – with high numbers of users and profiles, recertification of mainframe access can be tricky and time consuming. But this is a key control that safeguards your data.
- **Privileged access and accounts** – issues to consider include the number of accounts with privileged access, is this access recertified regularly, are privileged accounts behind a break-glass process, and is the usage of such accounts logged, monitored and alerted on?

Encryption: regulations, standards and best practice increasingly require the use of encryption. Risks that still need to be addressed include key management (where and how are keys generated? How are keys distributed? How often are keys changed/rotated? Are keys backed up?) and ICSF and cryptographic services (are the key datasets protected correctly? Are the datasets backed up and ae the backups protected? Is key usage audited? Who has access to the ICSF Panels?)

Dealing with complexity and ensuring observability

From conversations internally and with clients, we are seeing increasing awareness and interest around complexity and observability. These will continue to come to the fore, and of course are closely linked to wider issues around cyber security, how AI can be used, and why getting the basics right is so important.

So many mainframe deployments now sit at the centre of a highly complex web of applications and services. As we've seen, especially during and since the pandemic, the modern mainframe is now a hub for digital transformation. To help deal with this complexity, we will continue to see AIOps tools and approaches that combine AI and automation to streamline and optimize systems management and operations. Observability is critical, of applications and systems, and today that means extending our horizons further, across different software and servers, for multi-cloud, on-premises mainframe, and hybrid environments. Having that 360° view is more important than ever. Security matters feed into and out of all that, and AI has the potential to be a golden thread running throughout. Let's see what's changed, and what's new, when we gather our thoughts in 12 months and look ahead to 2026.

Leanne Wilson
Senior Technical Delivery Manager / Senior Security Consultant
Vertali Ltd.

With more than 13 years' experience in mainframes, systems engineering and cybersecurity, Leanne Wilson leads Vertali's mainframe technical delivery of security and infrastructure projects. She focuses on helping organizations around the world to secure, protect and optimize their mainframe infrastructure and related applications.

Follow the Leaders: 5 things we learned from the latest tech stack trends

Mike Dickson
Head of Product Marketing
Broadcom Mainframe Software

When it comes to business, getting your tech stack right is a big deal. The ideal mix of cloud, mainframe, and distributed servers can strategically slash IT costs and even outperform rival organizations. With IT now consuming over 34% of total operating costs, how businesses balance their tech stack can be the difference between leading the market or falling behind.

Last year, Broadcom partnered with Rubin Worldwide, the world's leading researcher in business and technology economics to examine how businesses are using their tech stacks and what it means for their bottom lines. What we found was compelling. Leading companies are investing more in cloud and mainframe to outpace the competition.

A lot has occurred over the past year, however. Markets have shifted. Trends have changed. GenAI has taken center stage. And consumer expectations have continued to rise. What does this mean when it comes to IT strategy and tech economics? To get answers, we tapped Rubin Worldwide again to analyze how the latest shifts and trends have impacted our inaugural "Technology Asset Class Optimization" study.

After crunching the numbers and surveying the most recent technology investment choices and performance of 2,400 global companies, this is what we found.

1

Striking the right balance reduces businesses’ cost per transaction. In a first-of-its-kind analysis of the “Technology Cost of Goods” - that is, how much it costs to complete one transaction - this year’s report found that companies who are running more workloads on cloud and mainframe are able, on average, to complete transactions for less cost. This means they’re getting a better return on their IT investment; more for less. This trend is true across industry sectors, but especially pronounced in transportation, retail, healthcare, and finance.

For example, best-in-class airlines who invest more heavily in mainframe reduce their cost per passenger by over 35% when compared to both cloud heavy companies and average performers (e.g. $7.63 cost per customer for mainframe heavy orgs vs $10.78 per customer for average performers and $11.71 per customer for cloud heavy orgs). Furthermore, mainframe heavy companies enjoy a 20% reduction in the cost per retail transaction, cost per hospital bed, and overall cost in insurance coverages when compared to cloud heavy and/or average performers. (See full chart for details).

2

Hybrid IT still dominates. And for good reason. Consistent with last year’s findings, a hybrid cloud and mainframe stack outperforms any individual asset class when it comes to improving lead time, change failure rates, and mean time to recovery (MTTR). It’s no wonder top-performing organizations earmark the bulk of their IT spending on cloud and mainframe. More specifically, best-in-class performers invest 10% more in cloud and 10% more in mainframe on average.

3

Inflation is cooling the cloud. With IT inflation running high, cloud is getting more expensive, not less, and its usage is still below that of mainframe and distributed for both average and top-performing organizations. While cloud is undoubtedly still a staple feature of many high-performing hybrid architectures, “cloud first” thinking is slowing and becoming more selective as companies look for “safe harbor” options, such as mainframe. This trend is reflected in a modest but continued downward shift in cloud spending (measured both as a percent of Global IT Services and as a Percent of Global IT Spending).

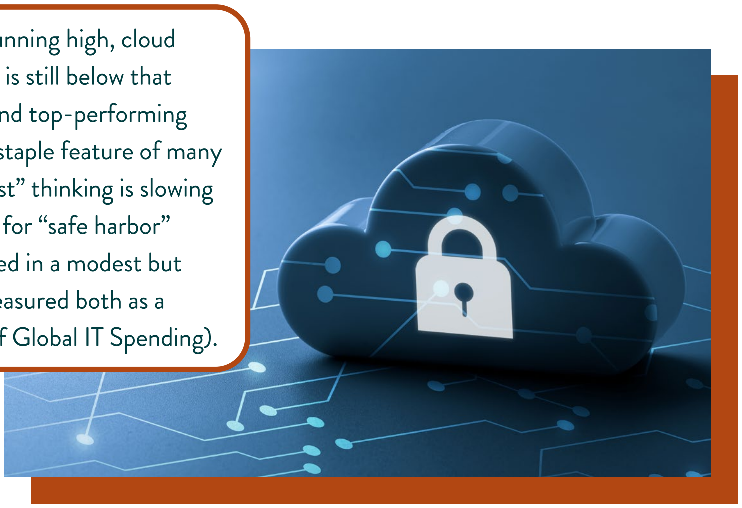

4

The value of Mainframe continues to increase. Since last year when the first edition of this study made abundantly clear that, “[n]othing scales like the mainframe,” the latest found even more efficiencies and performance gains at leading organizations who rely on mainframe. Over the past year, the unit cost savings of mainframe increased from 60% to 67%, whereas cloud and disturbed were slightly down at 20% and 10% to 18% and 9% respectively. The takeaway: hybrid organizations who lean in on mainframe are lowering their data, transaction, and overall IT costs of goods.

5

“On-premise” is still the majority. Similar to last year’s findings, 87% of global compute power happens on-premise. And as organizations rebalance their tech assets, this year’s study found that “some public cloud workloads are being repatriated to on-premise platforms.” That’s largely due to “unforeseen costs” of running certain workloads in the cloud and “unexpected expenses” to re-architect and migrate to the cloud.

In conclusion, this year’s Technology Asset Class Optimization report found that mainframes lower transaction costs of high volume environments demonstrably better than any other technology. On top of that, organizations that are “mainframe heavy” consistently outperform those with an overreliance on cloud or distributed. This is especially pronounced in travel, financial, retail, and insurance sectors. And it applies to established technologies as well as emerging ones—generative AI very much included.

The good news is we now have the data and models that allow IT leaders to evaluate and plan technology investments with the same rigor and language as financial investments. By examining their technology contracts against planned business volatilities, IT decision makers can consider the scalability of each technology asset before deciding which mix and balance is right for them.

While much has changed over the last year, one thing remains the same: Following the tech stack choices of leading organizations can have a measurable impact on your bottom line, and mainframe’s influence and value is playing an even greater role. Is your tech stack optimized to deliver the highest economic impact? To learn more, please get in touch or read the latest Technology Asset Class Optimization report.

What Steam Can Teach Mainframes

Allan Zander
CEO, *DataKinetics*

My father passed away recently, and I found myself thinking about him. He had many hobbies, but only one that caught my attention: building small-scale models. He had a steam train from when he was a boy, and at times, he would work on a small model railroad. Of course, I "helped" him, which sparked a bit of passion and interest in wanting to build a model railroad of my own.

Reflecting on his model steam train led me to think about steam as a transformative technology. Do we mainframe enthusiasts have a lesson to learn from the steam engine?

As a mainframer, I am quick to respond to the "all aboard!" conductor's song that the mainframe isn't going anywhere. We insiders know how much mainframes power the world's financial systems, how important the code is that runs those systems, and how critical the mainframe infrastructure is to those networks. It's true that ATM transactions, credit card processing, consumer data analytics, census number crunching, payroll processing, transaction recording, and insurance underwriting occur thanks to Big Iron.

But let's be cautious. While I believe there is comfort for many years ahead in the mainframe space, the one thing we can count on is change.

The First Act

When the steam engine was created, it was seen as a marvel. Sure, some critics noted that steam engines were remarkably expensive, and they required a team of talented individuals to keep them running and humming along. Sound familiar? However, the engineering and systems supporting steam energy were solid. The developers thought that they could control steam to an amazing extent.

It was not only the steam locomotive that changed the world; harnessing steam was a key catalyst of the Industrial Revolution. Steam appeared everywhere that needed mechanical energy, and people were convinced that the world would run on steam.

“Already the steam-engine works our mines, impels our ships, excavates our ports and our rivers, forges iron, fashions wood, grinds grain, spins and weaves our cloths, transports the heaviest burdens, etc. It appears that it must some day serve as a universal motor, and be substituted for animal power, waterfalls, and air currents.”— Sadi Carnot, 1824

Even though steam was life-changing, it didn’t stop with moving a locomotive. Next came superheaters created to heat steam beyond the temperature at which water boils. (Imagine convincing a room of executives that now you need extra special boiled water).

A steam superheater’s primary benefit is that it significantly increases the efficiency of a steam power plant. Superheated steam has a higher energy content than saturated (normal) steam, which allows for better turbine performance and improved heat transfer in industrial applications. It also reduced the risk of hydraulic water surges (water hammer) and corrosion.

Act Two: It’s Electric

Of course, progress dictates that change is inevitable. A second Industrial Revolution hit as electricity became cheaper to transfer to homes and businesses than steam. As might be expected, electricity faced critics too. Electric companies took out ads aimed at businesses, touting the benefits of electricity for efficiency and safety.

Electricity was considered more of a scientific curiosity than a useful phenomenon until the last part of the nineteenth century. I imagine the people behind steam saying things like, “Well, of course, electricity is an option, but that’s mostly replacing gas and for lights. Factories have large investments in steam power, and they’re not likely to change and adopt electrical engines and systems. I can see niche applications for electricity – maybe for locomotives that need to run underground where steam isn’t as practical, but electricity depends on a huge, distributed infrastructure. I think we will see steam for a long while.”

Thomas Edison opened the first commercial power plant, and now, electricity is ubiquitous. We hardly think about it. When was the last time an ad tried to convince you to try electricity?

The Third Act

Then, the mainframe came along, pioneering and propelling the third Industrial Revolution: the digital age. Problems that previously took a human mathematician 20 hours to solve took about 30 seconds for the ENIAC. Industries requiring robust computing capabilities quickly adopted mainframes.

Like steam and electricity, once they were embraced, new industries built up around them. The mainframe became not only a key business system but, at times, even a source of differentiation. Competitors offered scalability and reliability to do more—and faster—than any other company.

I bet the conversations in the board rooms about buying the first mainframe sounded eerily similar to the conversations about buying the first steam engine for a factory, and the first electric wiring for a hotel.

That steam superheater novelty of its day is comparable to in-memory technology in the mainframe. In-memory technology helps address effectively manage I/O in a mainframe. I/O contributes to the mainframe's legendary scalability, but it's also among the highest drivers of cost. Managing memory means managing computing efficiency and computational performance. Improved efficiency and performance then increase the value of the investments, ultimately attracting more revenue –exactly what management seeks. And why not? Everyone wants to reduce maintenance costs while simultaneously increasing productivity with the same investment.

The digital era prioritizes speed, scalability, and real-time access to information, relying on in-memory computing to enable new applications and architectures. In-memory technology has advanced dramatically from early mainframe systems—where memory was limited and optimized for large-scale batch processing--to systems where vast amounts of data can be processed in real time for faster and more flexible applications across industries.

The Fourth Act

Here we are now, somewhere at the beginning still of the fourth Industrial Revolution, where the biological and computational worlds may start to fuse together. Things like AI, robotics, 3D printing, and quantum computing may create applications that we can't even conceive of yet.

Perhaps learning a bit from the steam revolution, I've built my career and a great company in DataKinetics. (DataKinetics is a bit like a specialized superheater company of the steam era, and we have risen to become the gorilla in our space). I'm proud of our mainframe heritage and still enamored by steam. I appreciate it as a mechanical engineer; it appeals to the whimsy of a one-day retirement project, and it brought me closer to my father as we bonded over one of his hobbies.

"But looking around today – I don't see many steam engines. One day that will be true for the mainframe."

In the meantime, like the steam engine, let's be grateful for a great platform. Let's innovate around it. Both the mainframe and steam boosted major industrial revolutions. I am proud to say – and I believe --that the mainframe will be around for a very long time still. However, complacency breeds quickly, and in being dismissive or singularly focused, the next thing you know, there are suddenly no steam engines.

The steam era and digital era are both marked by technological revolutions that reshaped societies, economies, and cultures. Both have driven massive societal change, but the digital era is arguably more profound in its speed and scope of transformation. Let's take pride in the platform that still powers the world's financial systems, but not have so much hubris that we ignore the history of steam.

All aboard!