K2’s Tech Update

Latest Accounting Technology Lab Podcast Episodes

Published(Website)(Apple Podcast)

- Amazon Music, Google podcasts
- AssurancePrep
- Aiwyn
- SOBI Analytics
- Microsoft 365 Plans and Options for Firms
- What is the Future of Desktop Accounting
- Stanford Tax
- Zoho Analyst Day

Recorded, Will Be Released Soon

- Sage Transform
- Aider
- Datamatics
- MakersHub.ai
- Digilence
- Ricoh scanners
- Canon scanners

Major Topics

- AI and machine learning trends
- Quantum computing in finance
- Understanding the growing impact of blockchain and digital currencies in accounting
- Software and hardware innovations

Learning Objectives

- Identify and assess the impact of AI and machine learning on accounting practices
- List potential applications and implications of quantum computing in finance
- Identify the evolving role of blockchain technology and cryptocurrencies in financial systems
- Specify examples of the practicality and benefits of using the latest software and hardware tools

A Few Recommendations Last Year

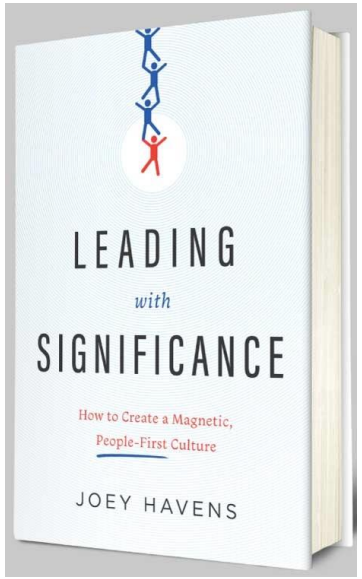

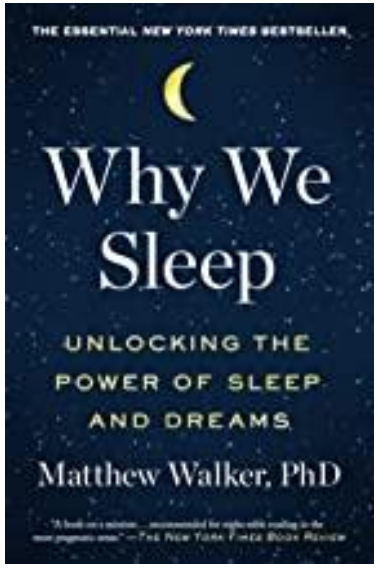

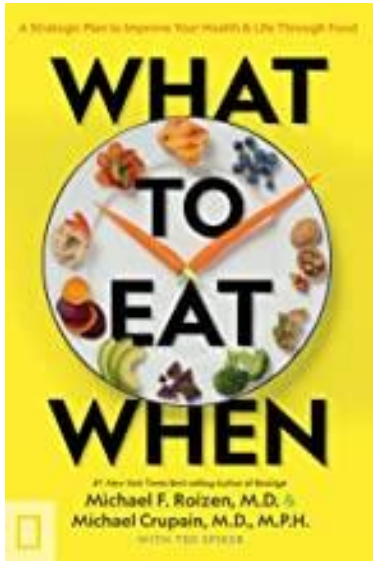

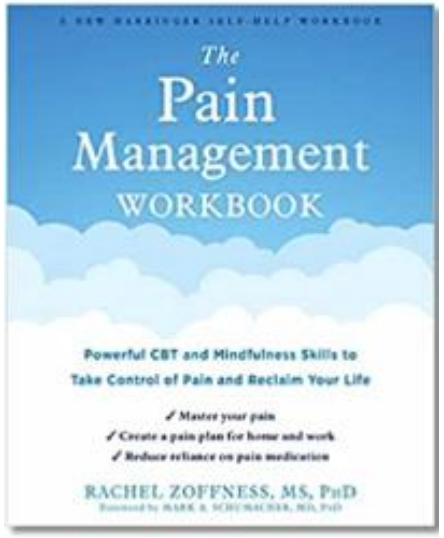

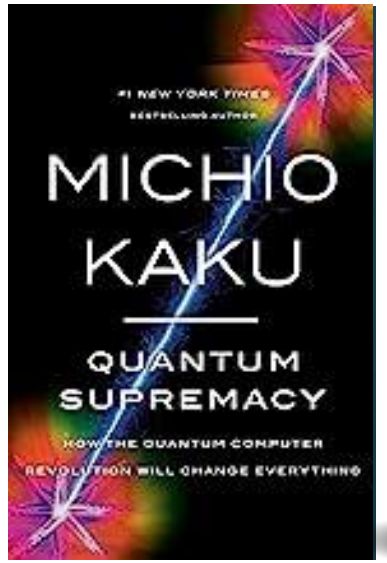

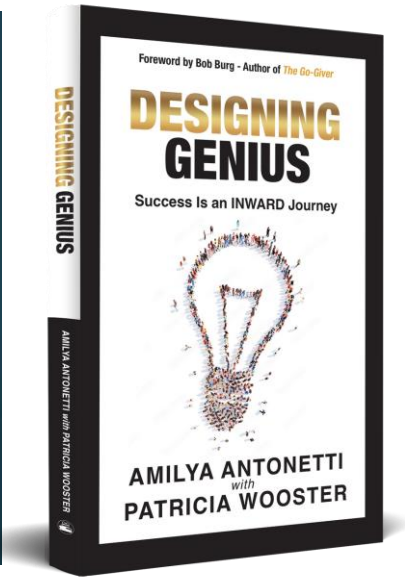

A Few Recommendations

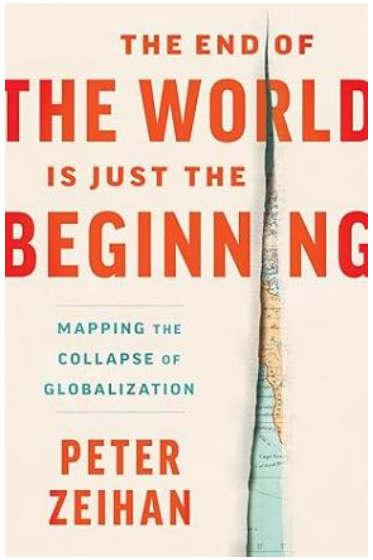

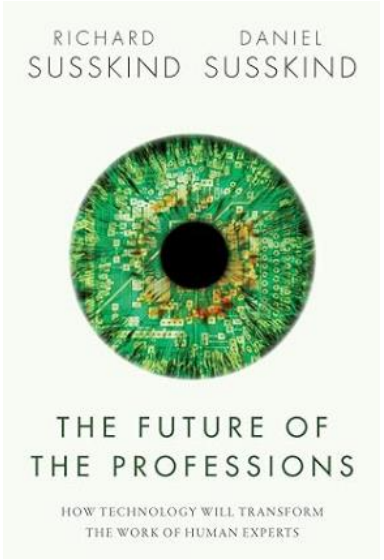

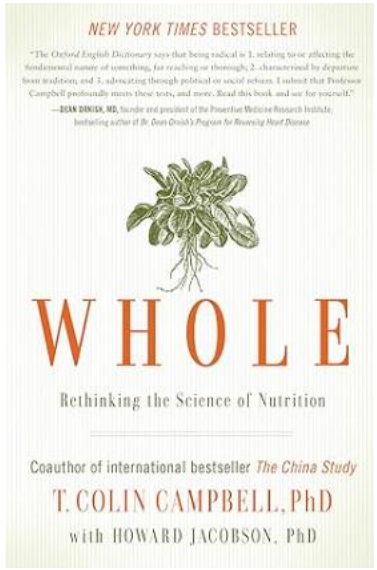

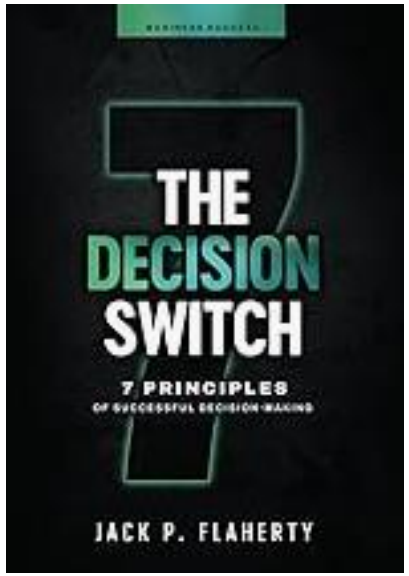

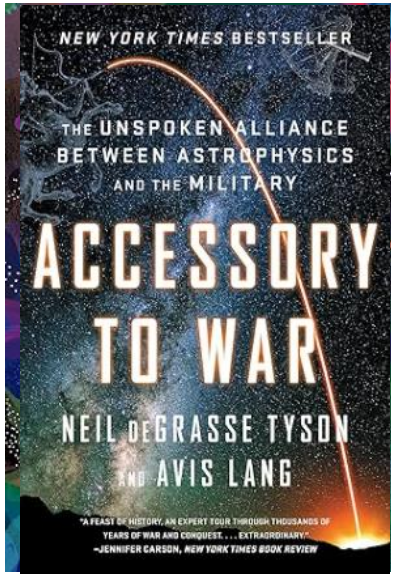

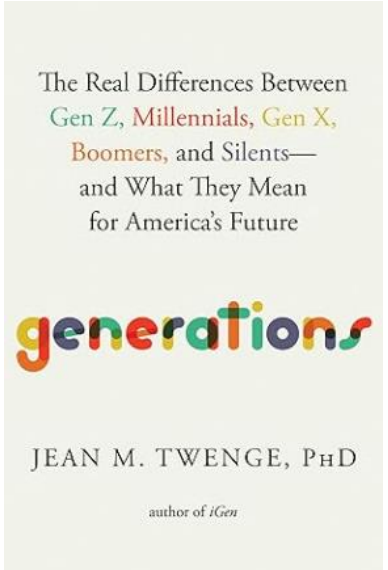

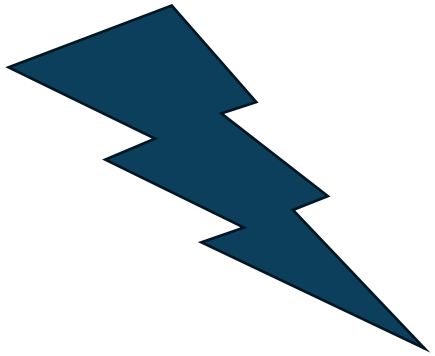

Last Year’s Super Buzz

ARTIFICIAL INTELLIGENCE

Artificial Intelligence Everywhere

- Almost every publisher has an AI strategy
- With Microsoft Copilot broadly available, Google's Gemini, and ChatGPT apps for everything else, productivity leverage is back
- Specific applications in accounting and finance are being released, and your stack is likely to get an AI upgrade
- Everyone wants to "be the expert"
- Courses abound everywhere, including from our K2 team
- Manulation = Manual + automation
- MVC chat AI connections with Open AI app store – Ben Thompson
- AI = Alien Intelligence or Apple Intelligence?

What Do You Need From AI?

Needs & Solutions

- Projections and forecasts?
- Productivity gains?
- Accurate information?
- Competitive advantage?
- Confidential information?

Enablement

- Policies
- Standardized procedures
- Creative use
- Sandbox development
- Protection

AI Leverage For Many

- AI acceleration
  - Large Language Models (LLM), but also Medium LM, Small LM, and Narrow LM
  - Models continuously refined
  - If we don't work, will we become like Wall-E?
- Thought Leaders
  - Stuart Russell, Hany Farid - UC Berkeley
  - Geoffrey Hinton (Nobel!), Eric Schmidt
  - Yoshua Bengio – Mila Quebec AI Institute
  - Yann LeCun - Meta
  - Alex Wang – Scale
  - Mustafa Suleyman, the co-founder of DeepMind, Inflection @Microsoft
- AI coverage in the media
  - *AI & The Future of Humanity*, The Whole Story, Anderson Cooper, 12/16/23, Nick Watt
    - Digital duplicate of Cooper by student Matyas Bohacek
    - NO FAKES Act discussion memo – US
    - Digital Services Act – requirement in the EU
  - *Artificial Intelligence: Its Promise…and Peril:* Fareed Zakaria GPS Special, 9/3/23
  - *The Daily Show: On The False Promises of AI:* Jon Stewart, 4/2/24
- Separate sessions on Artificial Intelligence

AI Considerations And Constraints

- Recursive Self-improvement, not under human control, 5 years
- Consider the impact of AI on jobs and careers of all kinds – the “tax of more people” according to Brian Chesky, CEO, Airbnb
- Tool use – learn something new and act on it
- Extreme risks – DeepMind – access to weapons
- Canonical AI – fulfill objective
- Wrong Objective Function – unintended and harmful behavior
- Rules = Constitution – Anthropic – ideas about limiting models
- Privacy regulation limiting inappropriate use of training data and data leakage should be expected
- Artificial General Intelligence – AGI, strong AI, full AI, human-level AI = sentience or consciousness – perhaps decades away

AI Effectiveness Comes From The Models, Courtesy Raju Vegesna, Zoho

ONE MODEL DOESN'T FIT ALL

There Is A Need For Multiple Models

- Narrow models
- Small Language Models (SLMs)
- Medium Language Models (MLMs)
- Large Language Models (LLMs)
- Large Behavior Models (LBMs) – for training robots

- That run on different platforms – Cloud, computer, phone

Narrow Models

- Purpose-driven, narrow models doing one task at a time
  - Grammar (text), data sets, NLP (Zoho's Ask Zia, Otter, Fireflies, Fathom)
  - Predictions (events like weather and self-driving cars)
  - Anomaly (time-series, trend analysis, etc.)
  - Email filtering
  - Voice assistants like Siri, Alexa, or Google Assistant
  - Recommendation systems like Netflix, Amazon, and Spotify
- Domain-specific
  - Finance, legal, security, etc.
- Fast and responsive, portability to mobile devices
- Multiple models deployed across Zoho
- Some vendors are trying to use a single model, frequently LLM

Small Language Models (SLMs)

- Models in the 3-7 billion parameter range
- Easier to fine-tune, less emergent (than MLMs and LLMs)
- Inference is done on the CPU (as opposed to GPU)
- Microsoft Phi-3 launched April 23 (10x cheaper than LLMs)
- Zoho's usage:
  - Translation
  - Advanced noise cancellation (Audio Video Stack)
  - Transcript generation

Difference In Model Recognition

- Narrow Model: $100
- Small Language Model: $9, which is correct and sufficient
- AI is improving legacy applications like OCR

Narrow Model: $100

Small Language Model: $9

Medium Language Models (MLMs)

- Models in the 20-50 billion parameter range
- Can be fine-tuned, more emergent than smaller models. Zoho is starting to deploy these in
  - Documents: Ask questions in your documents; mark anomalous documents
  - Audio Video: Generate intelligence over transcripts

Large Language Models (LLMs)

- Models in the 50+ billion parameter range
- Emergent behavior
- Not easy to fine-tune
- Hallucinations more common
- Zoho is integrated with public LLMs while investing in their own LLMs
  - Content generation
  - Find an anomalous legal document and recommend the corrected version

One Model Doesn't Fit All

- Narrow models
- Small Language Models (SLMs)
- Medium Language Models (MLMs)
- Large Language Models (LLMs)

We Expect Al Models To Be Commoditized

VALUE IS IN CONTEXTUALLY COMBINING MODELS

AI Applications

- Onward – Dave Marver, CEO, ARC-EX Brain Computer Interface (BCI), a competitor for Neuralink and Synchron BCI, which we covered in the past. Jocelyne Bloch, Neurosurgeon, Lausanne University Hospital, Gregoire Courtine, Neuroscientist implanted ARC-EX in bicycle accident patient Gert-Jan Oskam, Swiss Federal Institute Of Technology, Lausanne and can walk again
- Fake image of Pentagon bombing on a verified account that looked like Bloomberg News caused $500B in stock losses in two minutes – 5/22/23
- Tokamak Fusion controlled by DeepMind – Yves Martin, Dep. Director, Olivier Sauter, Sr. Scientist, Antoine Merle, Research Scientist, EPFL Swiss Plasma Center, Lausanne
- Waymo – Waymax simulator

AI In Art And Music

- AI Art
  - Dead End Gallery – Constance Brinkman
  - Unsupervised – MOMA
- Flawless AI for Cinematic Lip-Syncing video – Nick Lynes, Scott Mann
- Rules based music creation dates to 1950's
- Jukedeck by Ed Newton-Rex, 2012, is now irrelevant
- Native Instruments – RX corrects errors in things like pitch and timing inserting slivers of sound in "dynamic time-warping" which cleaned up Barbie and Oppenheimer soundtracks
- Yamaha – Vocaloid voice synthesizing tool translates and synthesizes singer's voice
- Suno – make a song about anything
- August 2023, Meta released the source code for AudioCraft now the basis for "infinite aspirations," promising composition according to Meta's Dr. Adi
  - MusicGen – 400,000 recordings, 28 months, 3.3bn parameters
  - Stability AI – Stable Audio, 800,000 tracks
  - Lifescore – recomposes original recordings, one client wants to release 6,000 AI versions of one track
- OpenAI – MuseNet
- Boomy – music automation platform, according to Alex Mitchell, CEO
- Voice-Swap – style transfers
- Artist Rights Alliance (ARA) letter **asks for ban on AI in music**. HR 6943 supported by human artistry campaign signed by 200+ top artists

AI Accounting Applications

Industry

- Botkeeper
- Digits
- MakersHub.ai
- Supervizor

Public Practice

- 4ImpactData
- Aider
- Black Ore Tax Autopilot
- Digilence
- laurel
- Wolters Kluwer TeamMate Document Linker
- Zoho

AI Applications Customized Yourself

- GPT4All - https://gpt4all.io/index.html
- Chatbox - https://chatboxai.app/
- Other third-party tools like Adept, Cohere, and Inflection
- Spreadsheet for prompt engineering
- Spend time learning how to customize Copilot, ChatGPT, Gemini or Claude
- Be aware of the rapidly evolving AI market with tools like Anthropic Claude, Baidu ERNIE, AI21 Jurassic-2, and Grok

AI As A Service (AIaaS)

- Makes advanced artificial intelligence solutions readily available to businesses through cloud computing

1. Google Cloud AI
2. Microsoft Azure AI
3. Amazon Web Services AI
4. IBM Watsonx
5. Oracle Cloud Infrastructure AI
6. DataRobot
7. OpenAI
8. Clarifai (vision)
9. BigML (machine learning)
10. H2O.ai (visualization)
11. RapidMiner
12. GroqCloud

AI Laws And Regulations

- Biden's US Executive Order 14110 - Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence – 10/30/23
- Bletchley Declaration G7 Agreement – 11/1/23, Hiroshima AI Process, Code of Conduct
- EU Artificial Intelligence Act – 12/8/23, 800 pages, currently highest bar
- CPFB – September (Credit Denials) and January (Adverse Action)
- US AI Safety Institute, established in early 2024 as part of National Institute of Standards and Technology (NIST)/US Dept of Commerce
- NIST AI Risk Management Framework (NIST AI 100.1) (January 2023)
- US General Accounting Office – "Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities" (June 2021)
- National Artificial Intelligence Initiative Act of 2020 (P.L. 116-283)

AI Privacy And Safety Concerns

- Nita Farahanty, Mental Privacy – Duke University
- Lt. Colonel Martijin Hadike, Dutch military robotics with armed THeMIS UGVs (unmanned ground vehicles) in Lithuania and cognitive warfare to gain access inside our brains using AI
- Jessica Dorsey – Utrect School of Law – legitimacy of military operations considering autonomy
- UN Secretary General wants a treaty to ban lethal autonomous weapons, but the US does not

The Next Big Thing Is Almost Here

QUANTUM COMPUTING

Quantum = 5th Industrial Revolution

- AI/Quantum acceleration
  - Impact on AI significant
  - Models run in parallel
  - “Nature is quantum from the start” – Sean Carroll
- Thought Leaders
  - Dario Gill – IBM
  - Dr. Charina Chou and Hartmut Neven (“Chief Optimist”) – Google
- Quantum hardware, encryption
  - 60 Minutes 12/3/2023
  - Michio Kaku Quantum Supremacy – “Quantum is the language of the universe”
  - Separate session on quantum

Examples Of Quantum

- Coin states – heads, tails and everything in between
- Maze – all possible paths at once

Quantum Computers & Encryption

- Y2Q (sometimes called Q-Day) – when Quantum Computers break encryption. Spy agencies and bad actors are warehousing data to be decrypted; that is quantum hacking
- Biden administration deadline for new encryption is 2035, but Shor's algorithm could break decryption by 2029 or sooner
- NIST has announced four quantum-resistant algorithms after a six-year competition: CRYSTALS-Kyber (FIPS 203, CRYSTALS-Dilithium (FIPS 204), SPHINCS+ (FIPS 205) and FALCON (digital signatures)
- Intel Tunnel Falls – 12 qubits, Google's Bristlecone has 72 qubits

Quantum Volume - Interference

- One measure of a quantum computer's capability is its number of quantum bits, or qubits
- But existing machines, which implement qubits in various ways, all have a fatal flaw: the delicate quantum states on which they depend "decohere" after a fraction of a second from interference
- A better measure may be so-called "quantum volume" (qv), which depends on the "width" of a computer (its number of qubits) and its "depth" (the number of operations they can perform before decohering). A computer with 14 qubits that can execute 14 operations is said to have a qv of 2 to the power of 14, or 16,384
- IBM has set a target of 2 to the power of 100

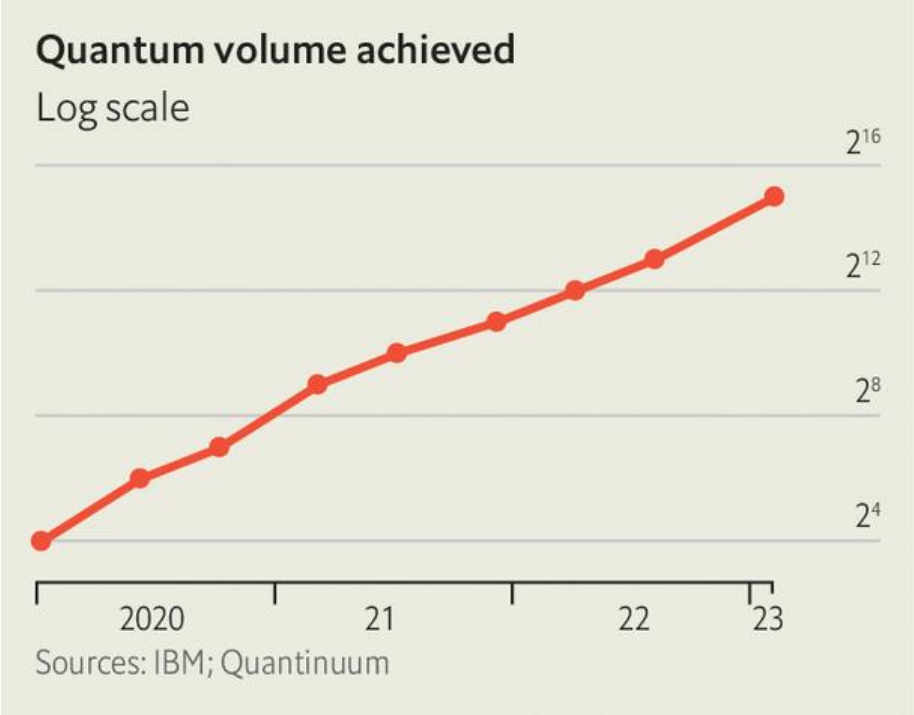

Speaking Of IBM Quantum…

- Condor, released on 12/4/23, has 1,121 qubits in a honeycomb pattern
- 433 in 2022, 127 in 2021
- Heron (pictured) has 133 qubits but a record-low error rate, three times lower than its previous quantum processor
- IBM System Two
- Error-correction techniques will require more than 1,000 physical qubits for each logical qubit
- Alternative error-correction scheme called quantum low-density parity check (qLDPC) promises to cut that number by a factor of 10 or more

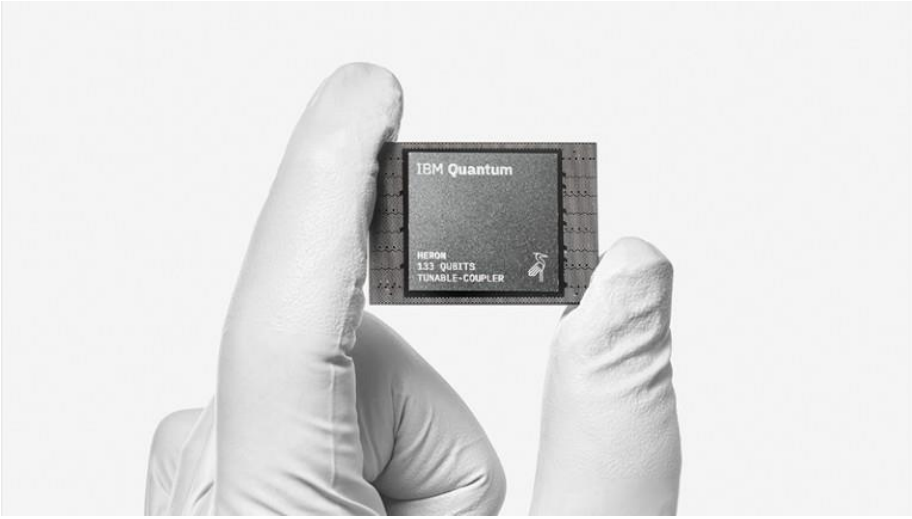

Microsoft And Quantinuum Errors

- Microsoft and Quantinuum say they've demonstrated a quantum computing system that can reduce the error rate for data processing by a factor of 800
- Qubits tend to be "noisy" — that is susceptible to perturbations that introduce errors
- They created four highly reliable logical qubits from just 30 physical qubits
- June 5 H2-1 56 qubits "three 9s"
- Will be available in private preview for Azure Quantum Elements

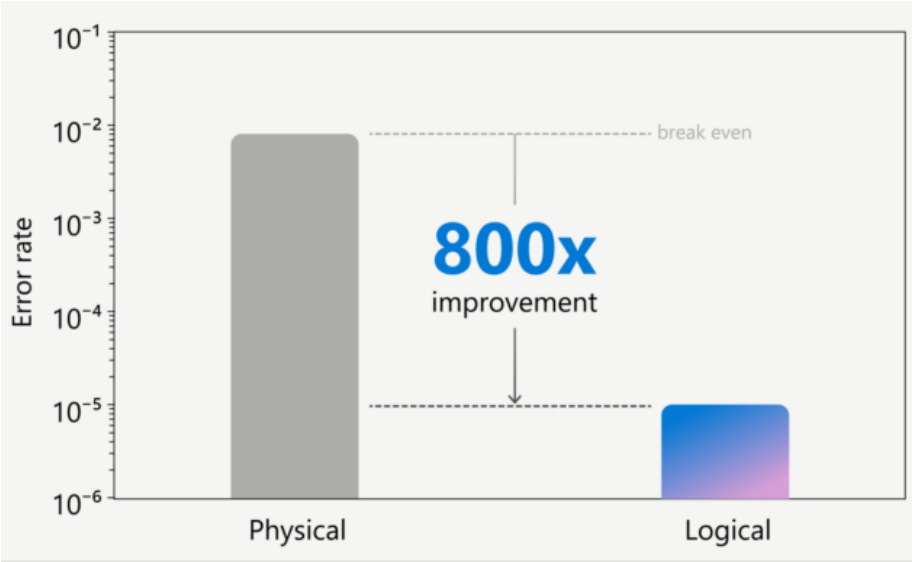

Quantum Entanglement

- LIGO gravitational wave detector utilizes entanglement within a laser beam
- Quark entanglement at the Large Hadron Collider reported on September 18, 2024
- Chinese Spy Satellites June 2017
  - The Micius satellite beamed entangled pairs of photons to three ground stations across China
  - These ground stations were separated by more than 1,200 kilometers (approximately 746 miles)
  - The test verified a mysterious and long-held tenet of quantum theory: entanglement remains intact even when particles are far apart

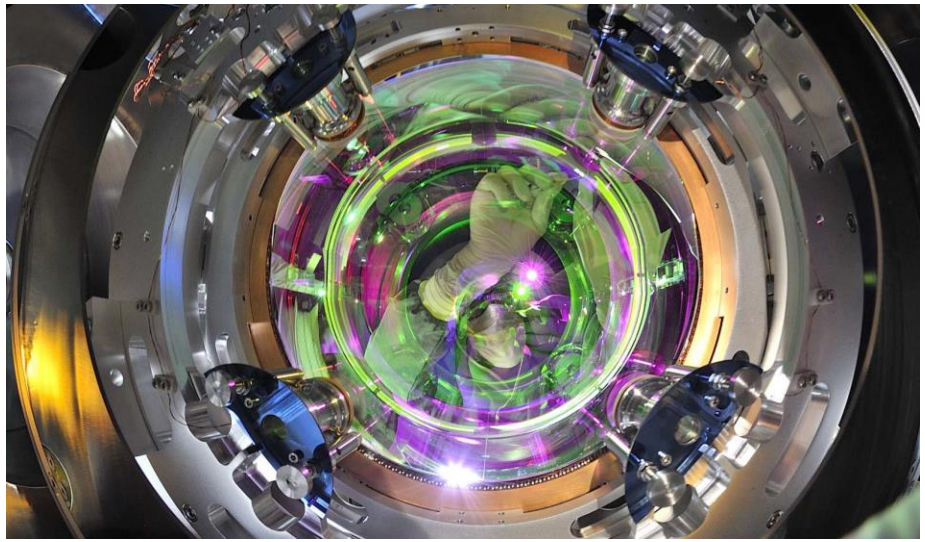

Quantum Applications

- Cleveland Clinic – Dr. Serpil Ezurum
- Model the behavior of proteins, which change shape to change function. When wrong, this causes disease
- Cancer and autoimmunity are the first targets

Quantum Applications

- Energy generation and storage
- Climate modeling and prediction
- Feeding the planet
- Drug discovery and materials science
- Gene editing and curing cancer
- Immortality
- Cryptography and security
- Optimization problems
- Fundamental scientific research

Quantum Accounting Applications

- AI Acceleration
- Optimization problems
- Fraud detection
- Cryptography

Quantum As A Service

- IBM Quantum: Provides superconducting quantum computing services
- Oxford Quantum Circuits (OQC)
- Rigetti Quantum Cloud Services: Offers superconducting modalities
- D-Wave Systems: Focuses on quantum annealing
- Amazon Braket: Offers superconducting, ion trap, and neutral atom modalities
- Xanadu Cloud: Provides photonic quantum devices
- And many more, including Strangeworks, Coldquanta, QC Ware, Intel, Atom Computing, and Zapata Computing

Quantum Benchmarking

- DARPA said the Quantum Benchmarking Initiative, or QBI, has stemmed from two of the agency's programs: Quantum Benchmarking and Underexplored Systems for Utility-Scale Quantum Computing
- "Our opening position is skepticism," stated Dr. Joe Altepeter, of DARPA. "We're pretty sure whatever you're doing is not going to work."
- According to DARPA, the QBI program will verify and validate approaches that enable revolutionary advances in design, engineering, testing, and evaluation in the domain of fault-tolerant quantum computing
- Research firm IDC is forecasting that customer spending for quantum computing will grow from $1.1 billion in 2022 to $7.6 billion in 2027
- According to a June report by quantum industry analyst firm IQT research, the worldwide market for quantum networks will be near $1.5 billion in 2027 and grow to more than $8 billion by 2031
- Quantum key distribution (QKD) will be the primary revenue driver, followed by a rise in networks that use emerging quantum repeaters to connect quantum computers and quantum sensor networks, the research firm predicts

Quantum Timeframe

- Don’t see an obstacle to 2030 availability, according to Dario Gill
- No fundamental breakthroughs needed - Hartmut Neven
- Machines are clearly running today, and capabilities are expanded with each new release
- Programs for some applications don’t have to be rewritten, they simply run orders of magnitude faster
- Small quantum machines can outperform some of the largest traditional computers today

Public Practice, Industry, And Productivity

ACCOUNTING APPLICATIONS

Accounting Software

Small Business

- Accounting Power (US only)
- AccountingSuite
- CYMA
- FreshBooks
- QuickBooks
- Sage 50 US/Canada
- Spire Systems
- Xero
- Zoho Books

Mid-Market

- Acumatica
- Dynamics 365
- Exact Globe
- NetSuite
- Odoo
- Sage Intacct
- SAP Business One
- Many more…

CPA Firm Software

Established

- Wolters Kluwer
- Thomson Reuters
- Intuit
- IRIS/Hg Capital
- Many others…

Innovative

- Additive.ai
- Black Ore Tax AutoPilot
- Liscio/HubSync/Digilence
- SmartVault/Suralink/Zoho
- With new products arriving…

ZOHO

- Privately Held
- Publicly Responsible
- First Bootstrapped Unicorn
- Homegrown Apps
- No Azure or AWS
- Manage Engine Network
- Privacy #1 – No Tags
- Single Suite
- Single Portal
- SOC 2 & GDPR
- Breakthrough product with Zoho Practice

28
Years

15K+
Employees

60+
Apps

150+
Countries

100M+
Users

750K+
Customers

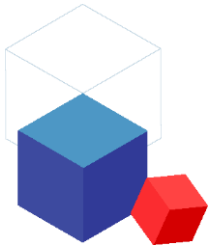

The Single Suite - The Future of SaaS

Home Automation Standards

- The Gadget Standard: **Matter**
- The Appliance Standard: **Home Connect**
- The Legacy Standard: **Z-Wave**

Home Automation Becomes Business Automation

- Conference room monitors - DeX/Bixby multitasks between devices including Intel, Apple, Microsoft, and Google all supported
- Voice recognition technology – Siri, Google, Windows, Alexa
- Break rooms - Qi v2.0 will unify the wireless charging industry under one global standard, including Ki2
- AR/VR headsets
- Computers for gaming
- Monitors for gaming
- NAS
- Tablets
- Transcription eyeglasses and Timekettle X1 Interpreter

Blockchain In Accounting

- Improved Transparency and Traceability
- Enhanced Security
- Increased Efficiency and Reduced Costs
- Smart Contracts
- Challenges and Considerations

Data Center Noise A Health Crisis?

- At least 137 datacenters in the U.S. run Bitcoin-mining operations 24x7
- Granbury, TX has 140 containers with 30,000 computers, producing 72 decibels and at times up to 91 = roughly between a lawn mower and a chain saw powering proof-of-work crypto mining
- German cardiologist exposed healthy students to just 63 decibels found vascular function diminished in one night
- TX penal code deems noise above 85 decibels unreasonable
- Federal mandate on noise was defunded during the Reagan administration
- Governor Greg Abbott stated that Texas would soon be the #1 state for blockchain & cryptocurrency. OK, IN, and MO urgently push Bitcoin Rights
- In March 2024 alone, Bitcoin mining generated $2B in revenue
- Data centers will use 8% of U.S. power by 2030, up from 3% in 2022

Data Center Noise Impacts

- Ear, Nose, and Throat specialist Salim Bhaloo sees new patients weekly with increased cortisol and sugar levels. “This thing is definitely causing a lot of stress. Everyone is just miserable about it.”
- Pulmonary embolism, atrioventricular block, hypertension, heart palpitations, chest pain, vertigo, tinnitus, migraines, panic attacks, blood pressure increase, loss of hearing, fluids leaking from ears, fainting, nausea, and depression
- Five-year-old Indigo Rosenkranz suffered from a seizure
- Animals going bald, shaking and hyperventilating
- Centuries old oak tree dies in three months
- Similar symptoms from people in AR and Williston, ND
- Are these reports a canary in the coal mine?

Datacenter Concentration Risk?

- After the July 19, 2024, global outage caused by CrowdStrike that affected 8.5 million Windows devices, IT teams are considering their risks
- Caused by a bug in Content Validator
- Delta Air Lines expects costs of $500M
- Fortune 500 total losses are estimated over $5B

- Alternatives:
  - Best overall CrowdStrike alternative: **SentinelOne Singularity**
  - Best for Microsoft ecosystems: **Microsoft Defender for Endpoint**
  - Best for transition into XDR: **Palo Alto Networks Cortex XDR**
  - Best for proactive endpoint protection: **Bitdefender GravityZone**
  - Best for SMB proactive endpoint protection: **Sophos XDR**

These Prototypes Show Promise As Emerging Technology

INTERESTING TECHNOLOGY TO WATCH

EVs

- Many products in the market from scooters to bicycles to motorcycles to cars to over the road trucks to off-road
- Charging infrastructure making progress, Tesla wins
- All of my family are only buying EV or hybrid
- Severe winter showed that better power sources needed
- Energy wind farms in Kansas worked when coal plants couldn’t produce energy

EV Charger Setup NMGI Office

Setup Considerations

- The amount of electricity you would spend on charging is a lot less than you would think
- Some companies let you set up a way to charge for charging or allow specific people to charge for free, but it can be a hassle
- Tesla chargers are what we picked due to their ease of use and ability to have both charging types in one charger without an adapter
- I understand that Tesla will let you limit the chargers to Tesla serial numbers you know if you do not meet their ‘business charging’ program requirements. I believe that is 8 chargers in one location
- One of the more popular places in our town for EVs, on their busiest month only, spent an extra 50-75 dollars in electricity for the people they offered free charging

H ?

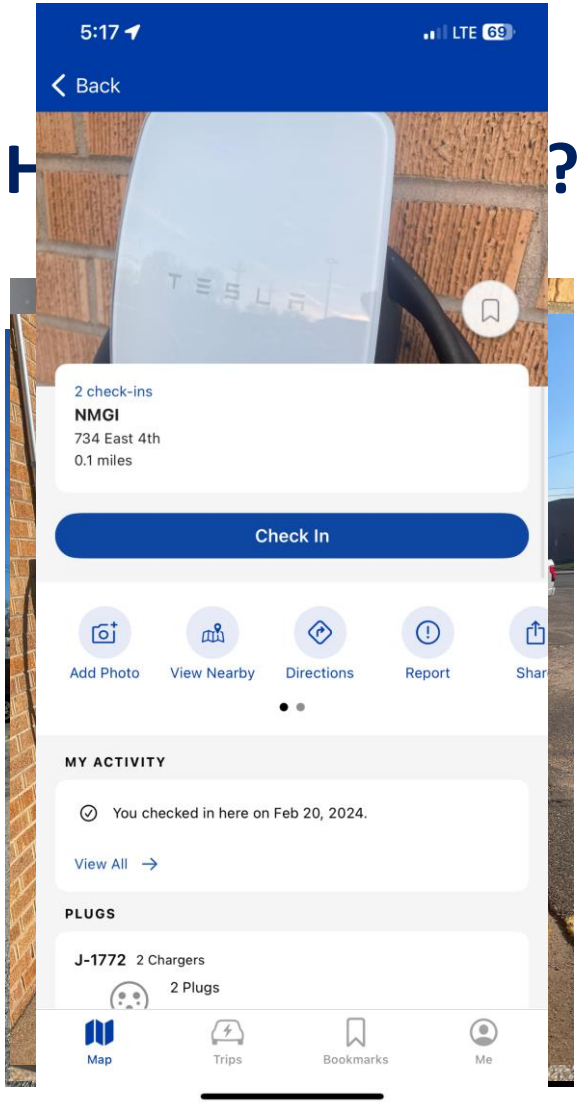

EV Charger Setup Home

Setup Considerations

- Most at-home chargers are anywhere from 16 to 50-amp chargers. Keep in mind that most electrical codes require you to have a breaker that is rated 125% above your charging rate. That means a 48-amp charger needs a 60-amp breaker
- Placement at home should be considered when putting in a charger
- It is probably best to put it toward your garage opening as some charge ports are in the back of the car, and some are in the front. If you put the charger in the front of the garage, you would never be able to reach your vehicle when it is parked outside
- Most chargers come with a 20–25-foot cable. Allow about 8-10 feet to be taken up by the cable lying on the ground

How Does It Look?

EV Charger Setup Home

Common Questions

- Tesla Chargers? $600 (now $620) + $3,000 new panel
- Ford Charger at home came with Lightning truck, about $600 used
- Time to charge? See chart →
- Only one vehicle on the Ford charger at home, mobile charger only adds about 3 miles per hour

What Type of Chargers?

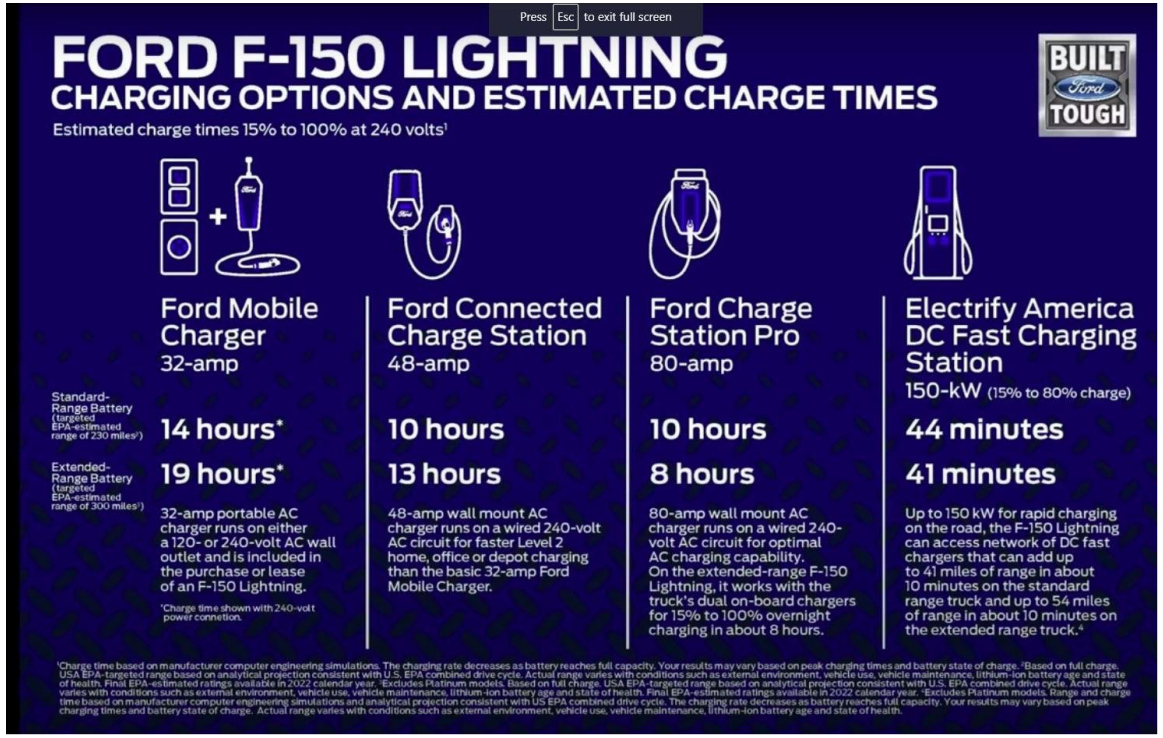

Detroit Testing Electric Vehicle Chargers Beneath Streets

- Demonstrations held November 29, 2023
- Copper inductive charging coils allow vehicles equipped with receivers to charge up their batteries while driving, idling or parking above the coils
- The quarter-mile segment of 14th Street will be used to test and perfect the technology ahead of making it available to the public within a few years, according to the Michigan Department of Transportation
- The technology belongs to Electreon, an Israel-based developer of wireless charging solutions for electric vehicles. The company has contracts for similar roadways in Israel, Sweden, Italy and Germany
- The pilot initiative in Michigan was announced in 2021 by Michigan Gov. Gretchen Whitmer

A Flying Car That Anyone Can Use Will Soon Go On Sale – No Pilot's License Needed For Helix

- Electrically powered vertical-take-off-and-landing (EVTOL) aircraft—flying cars, to the layman—are an idea whose time has not quite yet come but is fast approaching
- Helix is a single-seat vehicle, so "flying motorbike" is $190K
- 62MPH, range 18 miles, 4.5 hours recharge, 75 minutes on car EV chargers
- Joby Aviation flies EVTOL in New York City for the first time 11/13/23 and completed a 523-mile flight 07/11/24
- Consider Wisk autonomous

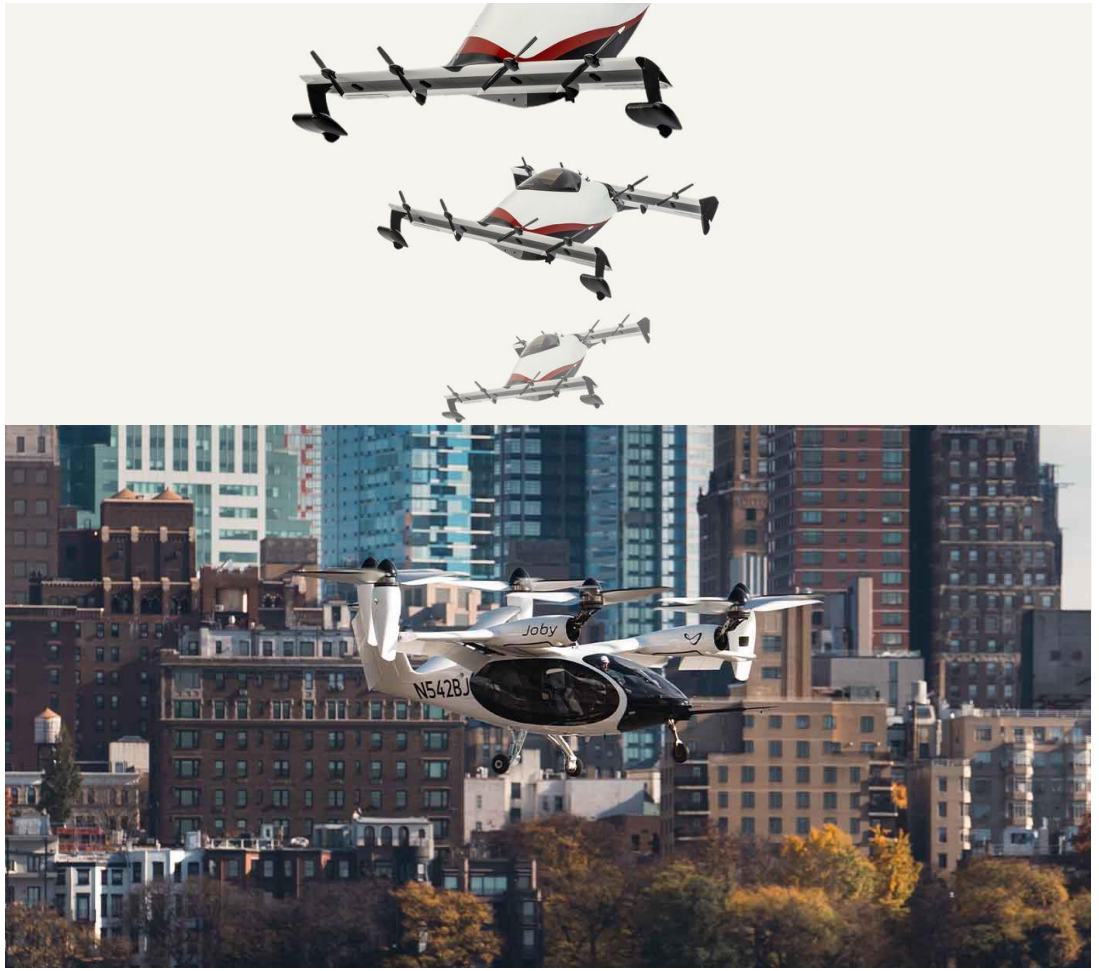

Or Perhaps Higher?

- SpinLaunch has a 1/3 scale prototype that doesn't require rocket propellent
- It is an electric-powered kinetic launch that is less expensive and sustainable using a vacuum-sealed centrifugal slingshot action
- Spaceport America in NM
- Successful launches on October 22, 2021, and in September 2022

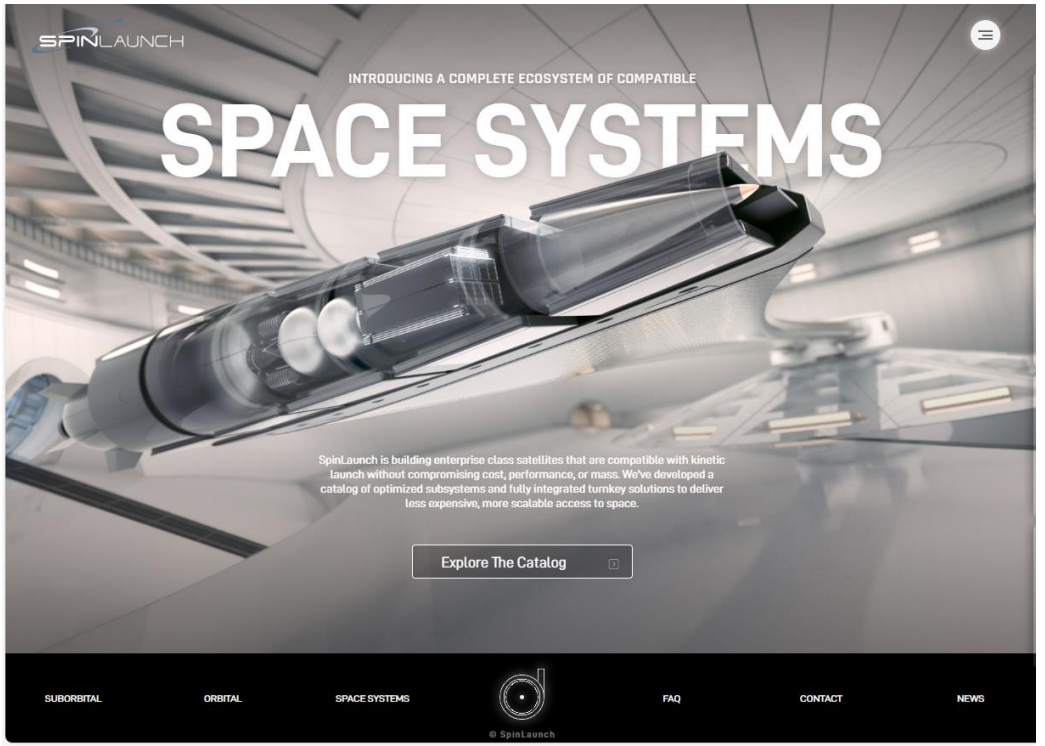

Or An Autonomous Boat For People Or Goods

- Roboat is a joint research program of Massachusetts Institute of Technology (MIT) and Amsterdam Institute for Advanced Metropolitan Solutions (AMS Institute)
- Tested in Amsterdam in 2017 & Boston
- Roboat uses Laser Image Detection and Ranging (LiDAR)

GNSS/PNT Refines 50-Year-Old GPS

- Vocabulary
  - GNSS – Global Navigation Satellite System
  - PNT – Positioning, Navigation and Timing
  - QZSS – Quasi-Zenith Satellite System
  - PNT – uses QZSS services to compliment GPS for more precision
- Benefits to GNSS/PNT
  - Vital systems
  - Robust security
  - Increased processing power
- Precision
  - Agriculture
  - Aviation
  - Facilities
  - Military
- GNSS correction
  - Precise Point Positioning (PPP)
    - Centimeter, 5 seconds
  - Real-time Kinematic Positioning (RTK)
    - 5-10 cm, 20-25 minutes
  - State Space Representation (SSR)
    - 5-10 cm, 10-30 seconds

Unfortunately, Surveillance Capitalism Is Expanding

- Skydio – AI for Drones
- BRINC – LEMUR2 has night vision cameras and can smash through windows purchased by NYPD
- Skydweller – doesn't have to land to recharge
- Axon Enterprise and Motorola for body cameras
- November 1, 2023, bill to prohibit buying drones from Chinese supplier DJI
- 200 firms selling space imagery, including SpaceX, private spy network, Starshield
- Blacksky can image every spot on earth every hour
- Remember Clearview from 2018?

Sodium-ion Batteries Are Coming Closer To Commercial Deployment

- Sodium is widely available, while lithium — presently the basis for most battery technology — can only be mined in a few places
- Sodium batteries have lagged behind on energy storage
- Swedish industrial startup Northvolt announced it had made a sodium battery that stored almost as much energy as a lithium one, aiming to reduce supply-chain dependence on China
- Meanwhile, China's own BYD has had similar ideas: It announced plans to build a multibillion-dollar sodium-battery factory, the world's first, hoping to sidestep the volatility in global lithium prices. It is not known how effective BYD's planned batteries will be

Hydrogen Batteries

- The world's most efficient system for producing hydrogen is scaling up for mass production. Some energy is always lost when turning water into hydrogen and oxygen — typical processes lose about 20-30%
- The Australian startup Hysata developed a new system, avoiding bubbles in the fluid that block the conduction of electrical charge and improving efficiency to 95%
- Hydrogen has advantages over batteries for long-term energy storage: It is more energy-dense — meaning it can be more easily used in industries such as aviation — and doesn't lose charge over long periods

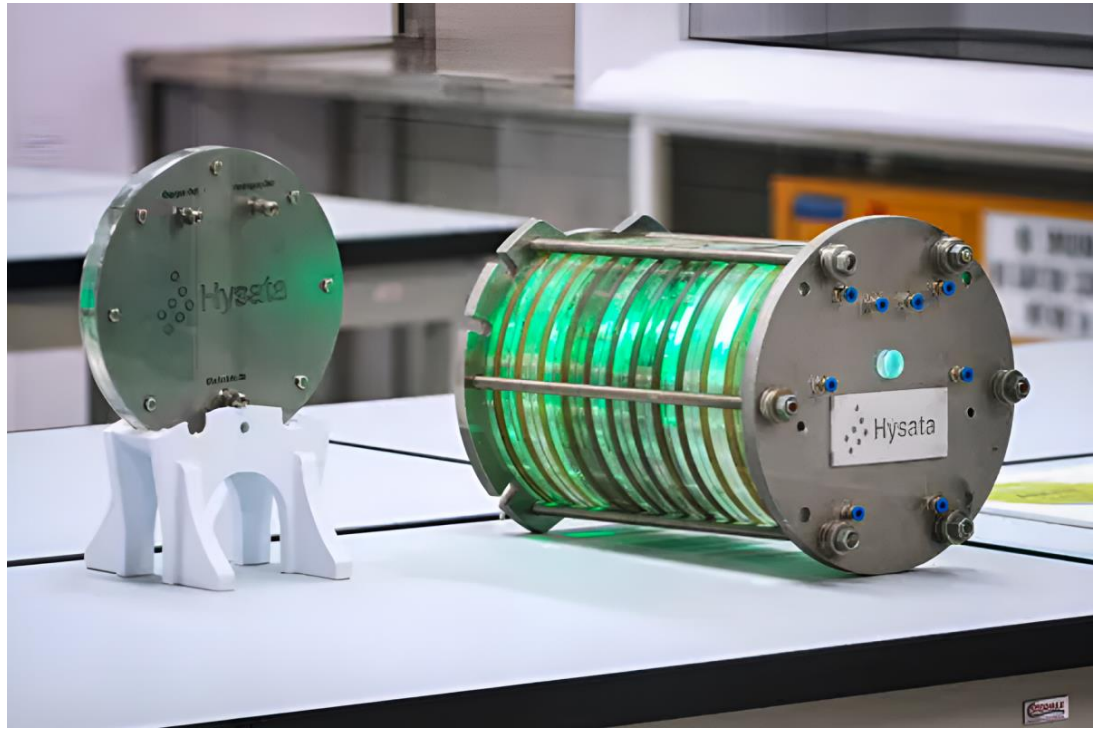

Nuclear Battery

- A Chinese company unveiled an “atomic battery” with a reported 50-year lifespan
- The Betavolt BV100 uses a radioactive nickel-63 isotope and new diamond semiconductor technology
- It does not provide a great deal of power, but could outlast the life of the tech it is powering: Its manufacturers envisage smartphones that never need charging or drones that can fly continuously
- Nickel-63 decays to harmless copper, and Betavolt says that unlike old nuclear batteries, which were dangerous and expensive, its newly developed product would not leak harmful radiation even if damaged or destroyed. The company hopes to release a 10-times-higher-output successor next year

In TVs, OLED Is Dead, QD-OLED

- Started with Samsung's S95B, Sony A95K, then the LG G2 with W-OLED
- Samsung S90C/S95C/S95D 55-83" $1,299-3,499
- MicroLED at CES 2024
- Watch for PHOLED in 2024 – phosphorescent OLED finally replaces the blue subpixel, reducing power use by 75%

PhoLED

- Watch for PHOLED in 2024 – phosphorescent OLED finally replaces the blue subpixel, reducing power use by 75%
- PHOLEDs utilize the principle of phosphorescence to achieve higher internal efficiencies compared to regular fluorescent OLEDs
- In all OLEDs, light emission occurs due to the electroluminescence of an organic semiconductor layer when an electric current passes through it
- Electrons and holes (charge carriers) are injected into this organic layer from the electrodes, forming excitons—bound states of electrons and holes
- Excitons can exist in either singlet or triplet states, depending on their spin properties
- While fluorescent OLEDs mainly emit light from singlet excitons, phosphorescent OLEDs generate light from both singlet and triplet excitons
- This unique property allows PHOLEDs to achieve nearly 100% internal quantum efficiency, surpassing the theoretical limit of 25% for fluorescent OLEDs

ATSC 3.0 – NextGen TV

- Built on the same Internet Protocol backbone as today's popular streaming media platforms
- Designed to bring together Over-the-Air with Over-the-Top content
- Delivers better video quality and immersive audio to viewers
- Consider https://www.freecast.com/connect
- Provides the capability of Advanced Emergency Alerting and Informing
- Easily adaptable to future technologies

Thunderbolt 5

- Doubles Thunderbolt 4 capacity to 80 gigabits
- 120 gigabits per second
- Supports USB4 V2
- Doubles PCI Express data

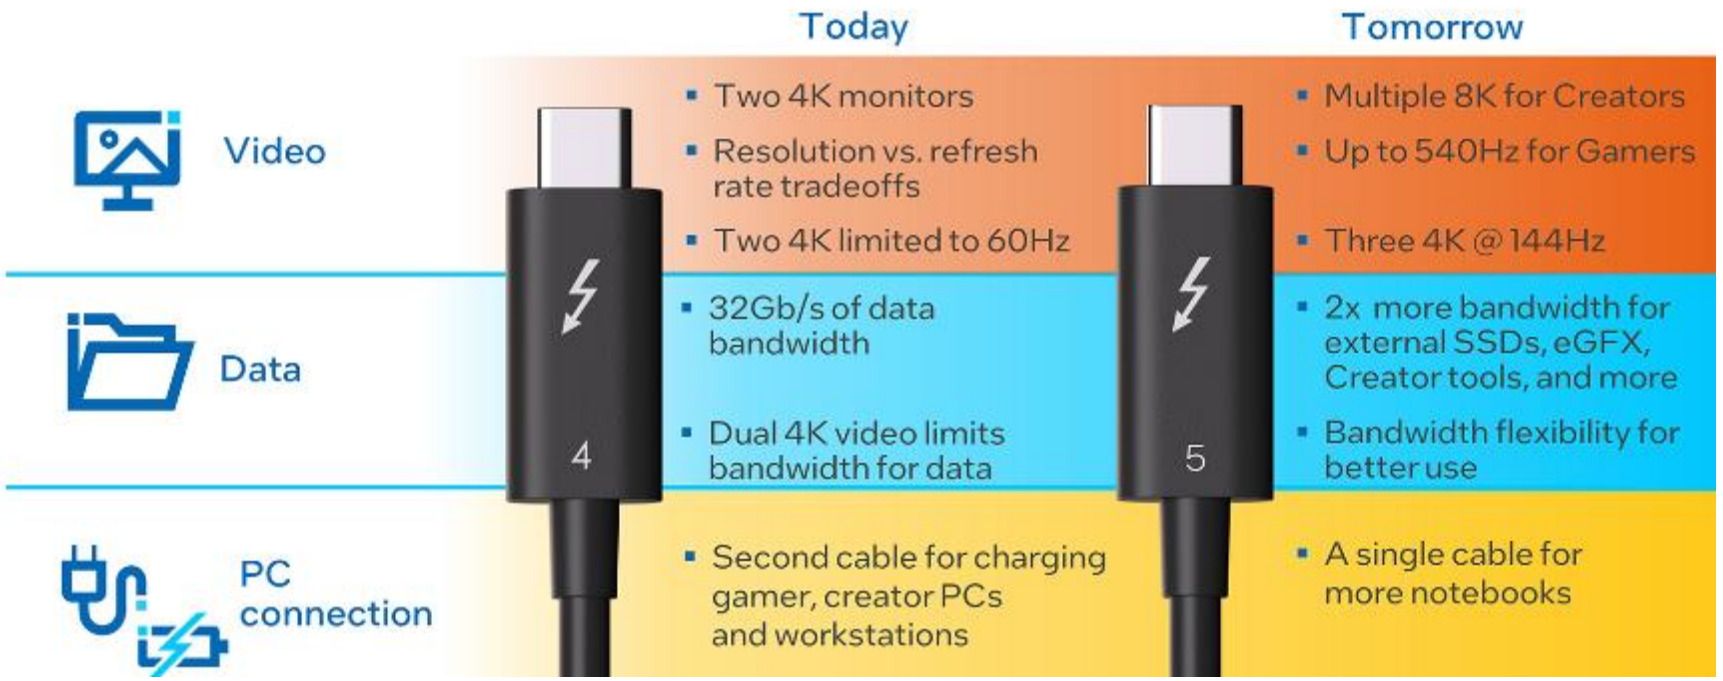

Types Of 3D Printers

Type of 3D Printer	Price	Speed	Complex	Precision	Main Characteristics
1. FDM (Fused Deposition Modeling)	L	M	L	M	Widely used, accessible, uses plastic filaments, popular for home use.
2. Resin 3D Printing (SLA, DLP, LCD)	L/M	M	M	H	Uses UV-sensitive resin, ideal for detailed models, delicate parts.
3. SLS (Selective Laser Sintering)	M/H	M	M	M/H	Uses polymer powders, suitable for industrial applications, no supports needed.
4. DMLS (Direct Metal Laser Sintering)	VH	M	H	H	Metal 3D printing, used in aerospace and automotive, expensive.
5. EBM (Electron Beam Melting)	VH	M/H	H	H	Similar to DMLS, uses an electron beam, suitable for industrial applications.
6. Binder Jetting (Ceramo One from CES was $5,000)	M/H	M/H	H	M	Uses powder and binding agent, capable of full-color models, versatile.
7. Multi Jet Fusion (MJF)	VH	H	VH	H	Advanced plastic 3D printing, extremely fast and accurate, industrial use.
8. PolyJet / Material Jetting	H	H	M/H	VH	Full-color and multi-material printing, used for detailed models, expensive.
9. Direct Energy Deposition (DED)	VH	VH	VH	M	Melts and fuses metals, used for large parts and repairs, less precise.
10. LOM (Laminated Object Manufacturing) / SDL	M	H	M	High	Layers of paper or plastic, full-color capability, wood-like parts.

Stacks, Recommendations, And Other Smart Options

PRACTICAL TECHNOLOGY CHOICES

Technology Strategies

- Adopt a productivity suite – Microsoft 365, Google Workspace, Zoho One (reverse backup)
- Attempt SaaS adoption where possible (mobile first, SOC Audit)
- Standardize your technology stack (fewer tools is better)
- Integrate and automate your processes and apps
- Assume continuous change (training plan)
- Leverage emerging technology (AI, SSO, Facial Recognition)
- Overall, the simplest approach has advantages

SSO – Pros And Cons

Okta

- Pros
  - Built-in contextual access for MFA
  - Customizable user portal
  - Multiple integrations with third-party applications
  - User self-service portal
- Cons
  - Can be buggy at times
  - May not work well with Microsoft apps

Microsoft Entra ID

- Pros
  - Integrates well with other Microsoft products
  - Geared towards enterprise use
  - Simple to understand subscriptions
- Cons
  - Contextual access to MFA only available on premium tier
  - No User SSO application

Other alternatives include Zoho Directory, JumpCloud, CyberArk, OneLogin, Passly (formerly AUTH Anvil)

Okta Pricing

Single Sign-on (SSO)	Adaptive SSO
$2 per user, per month	$5 per user, per month
• Okta Integration network. • Custom org URLs. • SIEM integration.	• All SSO features. • Location, device, network contextual access management. • Risk-based Authentication.

Multi-factor Authentication (MFA)	Adaptive MFA
$3 per user, per month	$6 per user, per month
• Security questions. • Okta Verify OTP. • Windows Hello and Apple TouchID.	• All MFA features. • Location, device, network contextual access management. • Network anonymizers.

Microsoft Entra ID Pricing

Entra ID Free	Entra ID P1	Entra ID P2	Entra ID Governance
Free	$6.00 per user, per month	$9 per user, per month	$7 per user, per month
• SSO unlimited. • Cloud authentication. • Customizable user sign-in page. • Partial MFA and conditional access.	• All Free features. • Group assignment to apps. • Cloud app discovery. • Full MFA and conditional access.	• All P2 features. • Identity protection. • Full end user self service.	Identity governance.

My Stack? Think About Yours!

Today

- Microsoft Teams/365/Adobe
- Grammarly/Cal.com/Bitwarden
- Copilot/Gemini/ChatGPT
- Defender/Sophos/CrowdStrike
- HP Dragonfly Gen4/ViewSonic
- ConnectWise for NMGI Management
- Zoho One for K2 Management
- iPhone/Apple Watch
- Stream Deck/Camtasia

Retiring

- Zoom/GoToWebinar
- Calendly/LastPass/YubiKey
- Cortana
- Webroot
- iPad
- QuickBooks
- ShareFile
- Citrix?

Well-Established, Replacements

Mainstream

- Wi-Fi 6E
- 5G
- Graphics Processing Units
- 4K
- 1 Gigabit Ethernet
- Smartphones & watches

What's Next

- Wi-Fi 7
- 6G
- Neural Processing Units (AI)
- 8K
- 100+ Gigabit Ethernet
- Wearable and AR/VR headsets

Practical Recommendations For Purchases

HARDWARE TRENDS

Fundamental Computer Technology

PC Buyer	Mac Buyer
• Windows 11, Microsoft 365	• MacOS, Microsoft 365
• Intel Core Ultra & AMD Ryzen 5-9	• M3/M4 PRO/MAX ARM
• Dedicated GPU (ARC, Radeon)	• Dedicated ProRes graphics
• 32-64GB of RAM, DDR5	• 8-64GB of RAM (DDR5 Unified)
• PCIe Gen 4 128GB-1TB of NVM Express Solid-State drive (NVMe)	• PCIe Gen 4 128GB-2TB of NVM Express Solid-State drive (NVMe)
• Thunderbolt 5, for HDMI/USB	• Thunderbolt 5, for HDMI/USB

Silicon Returns To Silicon Valley

- Chips produced in 1971 had 200 transistors per square millimeter. In 2023, the MI300, a processor built by AMD, crammed 150m transistors into the same area. The **smaller the transistors got, the faster they could switch on and off**
- In August 2022, the US government dangled a $50bn package of subsidies and tax credits to lure chip manufacturing back to America (the CHIPS Act). Other regions have followed suit, with the EU, Japan, and South Korea promising almost $94bn in handouts
- For much of the transistor's history, it followed a "happy scaling" path—as the logic gates shrank, they got faster and used less power. **That era is over**.

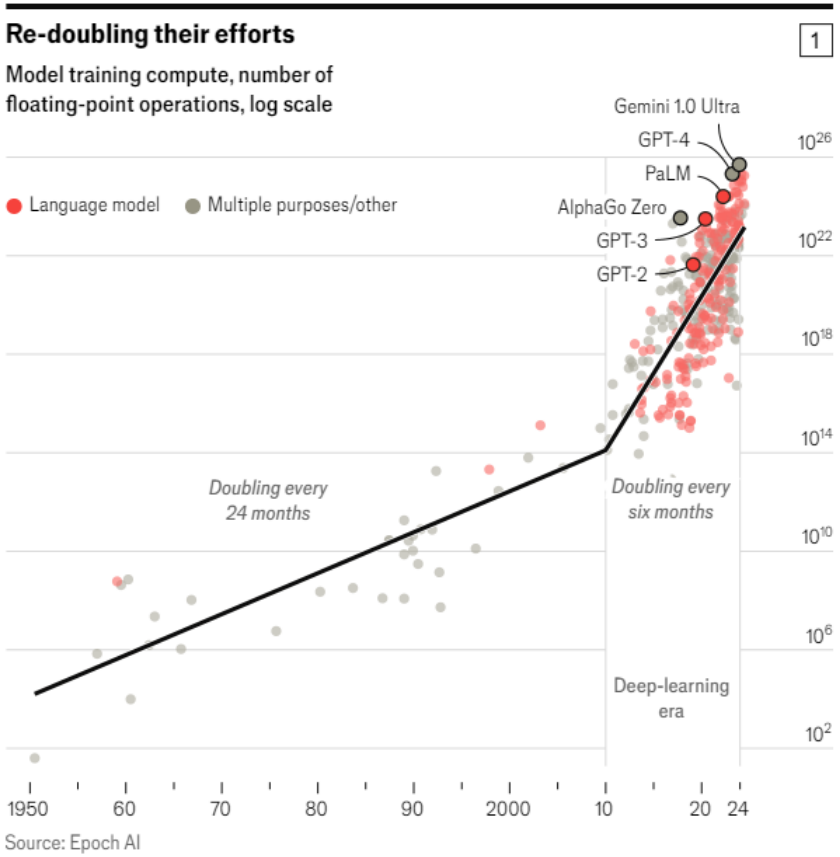

The Economist, Technology Quarterly, Sept 21-27, 2024

Don’t Believe The Label

- In the mid-1990s, **gate length started shrinking much faster than the half pitch**. Then, in the 2000s, problems with power and waste heat saw progress in shrinking gate lengths slow down sharply
- The measurements referred to by the companies caught up with the gate lengths and moved on down below them (see chart)
- At one point, Intel made chips in which the metal half-pitch was 40nm, the gate length was 26nm, and the fins within the finFET transistors were 8nm wide. They were referred to, somewhat arbitrarily, as “22nm.”
- **Unfortunately, 5, 4, 3, and 2 nm don’t mean what 22, 16, and 14 did**

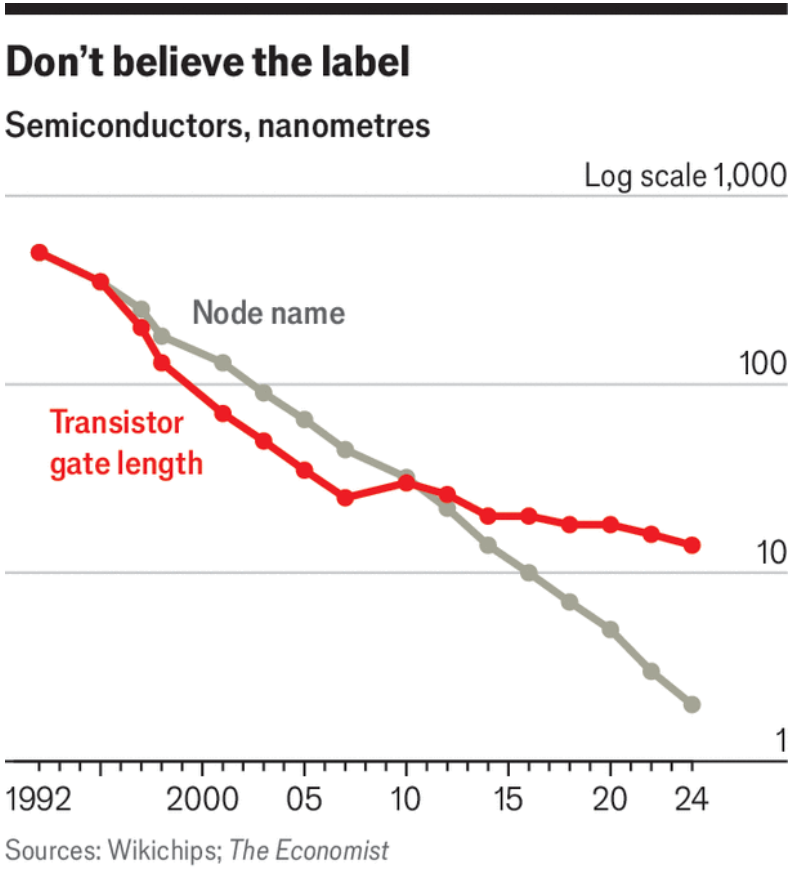

The Economist, Technology Quarterly, Sept 21-27, 2024

CPU, GPU, TPU/NPU

- **A central processing unit (CPU)** is an “everything machine” that can run any software, from word processing and email to spreadsheets and browsers. However, CPUs are not designed to meet neural networks' large-scale parallel processing needs. For every operation, they **sequentially load one data item** from memory, process it in the core, and store the results back in memory
- **A graphics processing unit (GPU),** the earliest example of this approach, has thousands of specialized cores for arithmetic operations crucial for neural networks. GPUs can **outperform CPUs by orders of magnitude** by focusing on specific tasks like multiplication and addition. Nvidia’s GPUs, designed initially to speedily render video-game images by processing countless pixels simultaneously, turned out to be very well suited to running neural networks, too
- **The tensor processing unit (TPU)** contains thousands of multiply-and-add units directly connected to a giant grid. The TPU **loads the data from external memory** into its grid, where it **flows through in regular waves**, similar to how a heart pumps blood. After each multiplication, the results are passed to the next unit. By reusing data from previous steps, the TPU reduces the need to access the off-chip memory

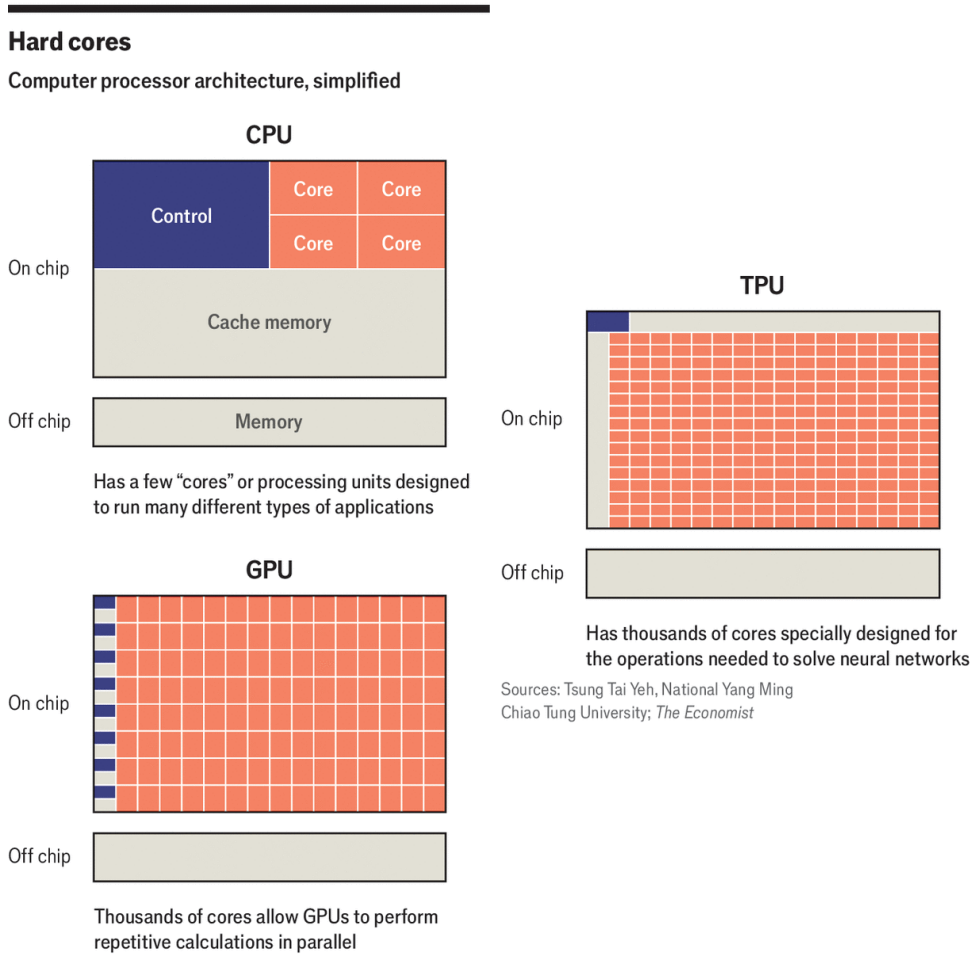

The Economist, Technology Quarterly, Sept 21-27, 2024

PC Sales Growth

- CPU shipments surpass levels during the pandemic era — demand for AMD and Intel processors has recovered
- Since Q4 2022, CPU shipments have increased 7% every quarter, for a 22 % increase
- iGPU (integrated graphics) are up 18% YOY, to 60M units

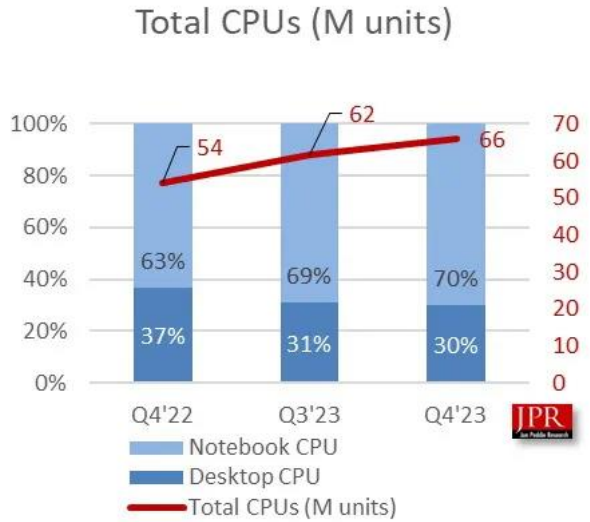

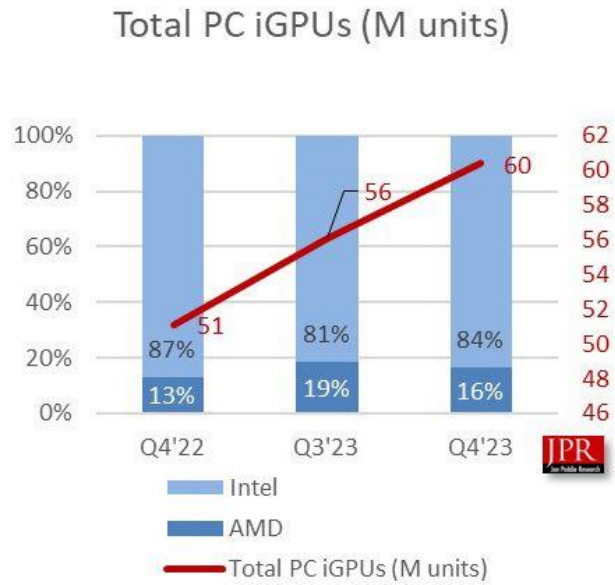

PC Sales Growth

- Up 8% to 267 million units
- Driven by new AI software
- AI-powered PCs will account for 19% of PCs shipped in 2024
  - M-series Macs
  - Ryzen 8000
  - Intel Core Ultra (Meteor Lake)
- Snapdragon X Elite could help ARM-based PCs grab 30% of the market share by 2026

Computer Chips

- Intel
  - Core Ultra (Meteor Lake, 14th Gen, “I” naming is gone)
    - Ultra = Neural Processor Units (NPUs)
    - AI integration
    - DDR-5 support
    - Better integrated graphics
  - Received $20B from the US CHIPS and Science Act for the AZ Intel Ocotillo Campus, OH, NM, and OR, 30,000 jobs
- AMD
  - Nine new Ryzen 8040 mobile with AI software for the seven with a Neural Processing Unit
  - Ryzen 8700G with integrated Radeon Graphics can almost match Nvidia’s GTX 1650, an entry-level gaming GPU

Computer Chips

- Apple M3, M3 Pro, M3 Max
  - 3nm
  - 2.5x the M1 rendering
  - All-New GPU Features Dynamic Caching, Mesh Shading, and Hardware-Accelerated Ray Tracing

Apple M3 chip specs: M3 vs M3 Pro vs M3 Max

	M3	M3 Pro	M3 Max
CPU	8 cores	12 cores	16 cores
GPU	10 cores	18 cores	40 cores
Memory	Up to 24GB	Up to 36GB	Up to 128GB

- Snapdragon X Elite – mid 24
  - Arm CPU faster than AMD's top-end mobile APU
  - 12x Intel Core Ultra 7
  - 12 Oryon 4.3Ghz cores (Nuvia)
  - TSMC 4nm

Snapdragon X Elite Performance in Geekbench 6

Windows

Sample chip	Single-Core	Multi-Core
Snapdragon X Elite	2,574	12,562
Ryzen 9 7940HS	2,475	11,667
Ryzen 9 7945HX	2,727	15,623
Ryzen 5 7600	2,736	12,260

AMD Claims Ryzen 79% Faster Than Intel Meteor Lake In AI Benchmarks

- Local AI processing in Llama 2 and Mistral Instruct 7B seem much faster on AMD
- Llama 2 is a state-of-the-art LLM from Meta, and Mistral Instruct 7B is an LLM with 7.3 billion parameters developed by ex-Meta and DeepMind developers
- ‘Time to first token'—which charts the time from inputting a prompt and pressing enter to the first results showing up
- AMD says that three kinds of cores can be used in concert: NPU, RDNA 3 GPU, and Zen 4 CPU
- Intel says Arrow Lake and Lunar Lake will have three times more AI performance for GPU and NPU. Intel's next-gen laptop and desktop processors are due later this year

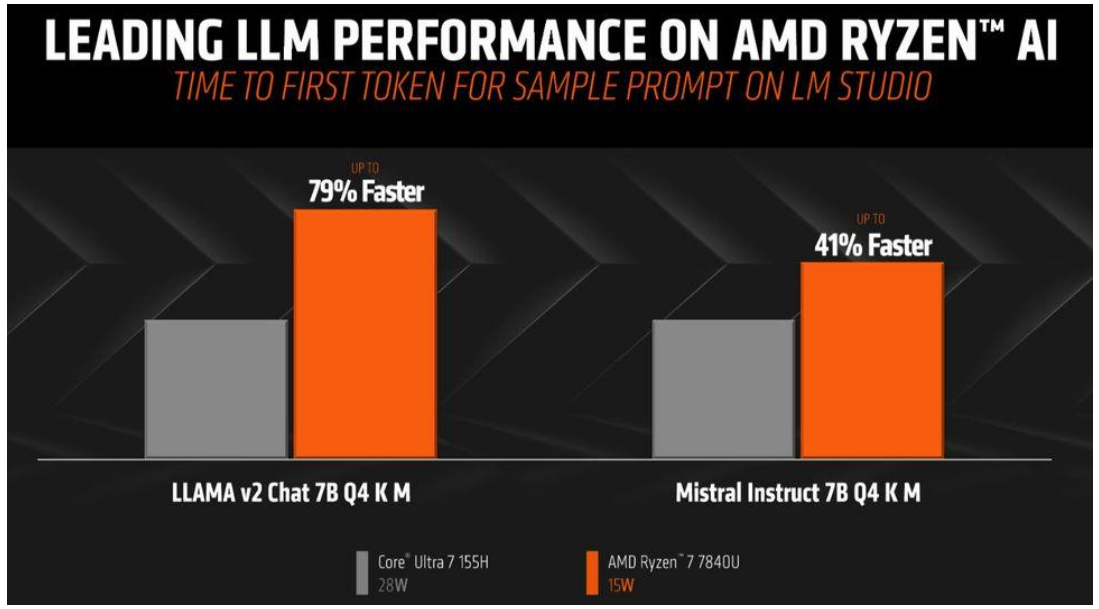

Phone Chips

A17 Pro/A16 Bionic

- 64-bit ARM-based system on a chip (SoC) built on 3 nm process by TSMC
- Apple-designed 64-bit six-core CPU with 8 GB of RAM:
  - Two high-performance cores run at 3.78 GHz.
  - Four energy-efficient cores run at 2.11 GHz.
- 16 Core Neural Engine capable of **35 trillion operations per second**
- Support for AV1 decoding and USB 3.2 Gen 2 (up to 10 Gb/s)
- A18 in the fall with the iPhone 16 (later)

ARM/Qualcomm Snapdragon 8

- Snapdragon 8 Gen 3, announced in October 2023, has 8 cores, 8 threads, Hexagon NPU with 34 TOPS
- X Elite has 12 cores, integrated Oryon CPU has 2x the power of the A17, GenAI models of 13B parameters, 4.5 the speed with 45 TOPS
- CPU Configuration: It has a 1+5+2 CPU setup, including one Cortex-X4 core, five Cortex-A720 cores, and two Cortex-A520 cores
- 30% faster than its predecessor while being 20% more efficient
- AI Capabilities: Designed for on-device generative AI tasks, it supports text-to-image generation, photo expansion, and large language models with over 10 billion parameters
- Used in Samsung Galaxy S24

AI Chips Vs. GPU

- Nvidia H-100, H-200, B-200
- Amazon – Inferentia/Graviton4/ Trainium2
- Cerebras – Giant chip, 900,000 cores, 44GB Memory, with 7,000 times the memory bandwidth of a GPU
- Google Tensor Processing Units (2015), Cypress/Maple/Axion
- Intel Gaudi 3
- Microsoft Maia 100/Cobalt 100
- Meta – Meta Training and Inference Accelerator (MTIA) & Research SuperCluster (RSC) and Artemis
- GroqChip – Language Processor Unit (LPU), clusters for memory access, 230 MB, with 80TB of memory bandwidth for 16 chip-to-chip interconnects
- 7 of the 10 most valuable companies make chips

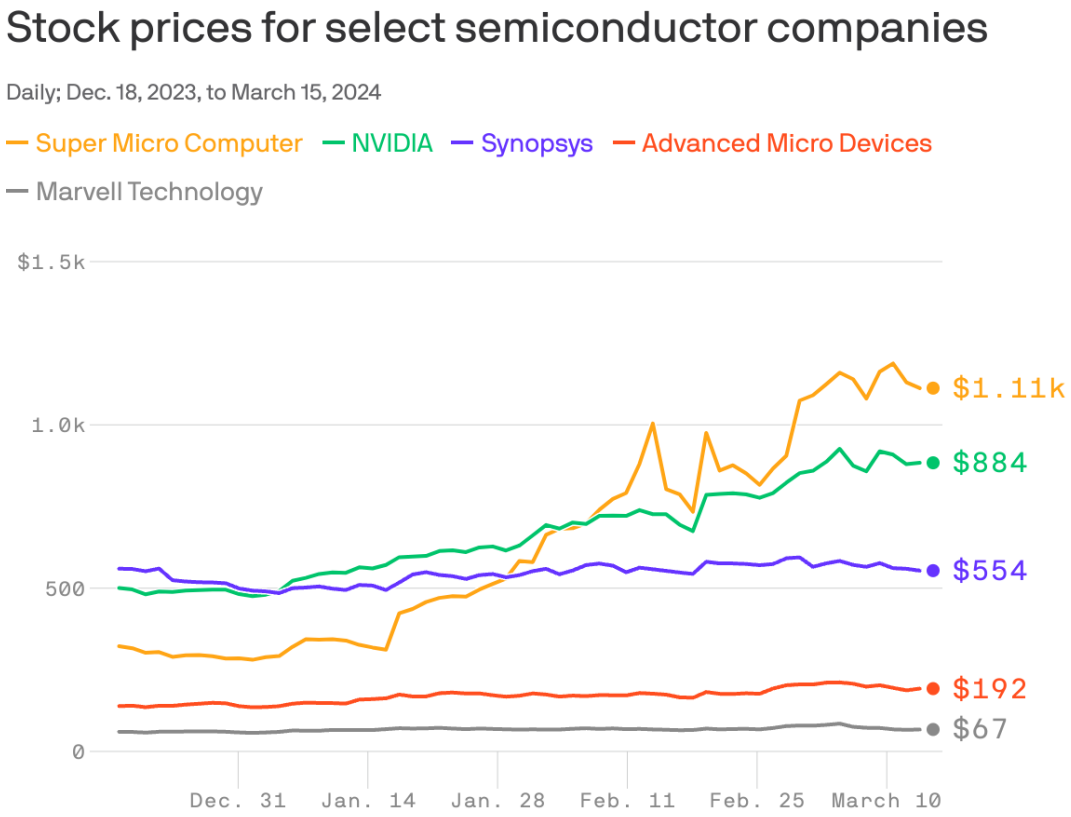

Nvidia Blackwell B200 And GB200

- Nvidia CEO Jensen Huang holds B200 on the left and H100 on the right
- 20 petaflops of FP4 horsepower from its 208 billion transistors
- GB200 combines two of those GPUs with a single Grace CPU that can offer 30 times the performance for LLM inference workloads
- “reduces cost and energy consumption by up to 25x” over an H100 according to Nvidia
- Training a 1.8-trillion parameter AI model takes 8,000 processors and 15 megawatts on existing tech, but 2,000 Blackwells would use just four megawatts for the same result
- **Rubin AI in 2026 using HBM4 (high-bandwidth memory) and NVLink 6 Switch at 3,600 GBps**

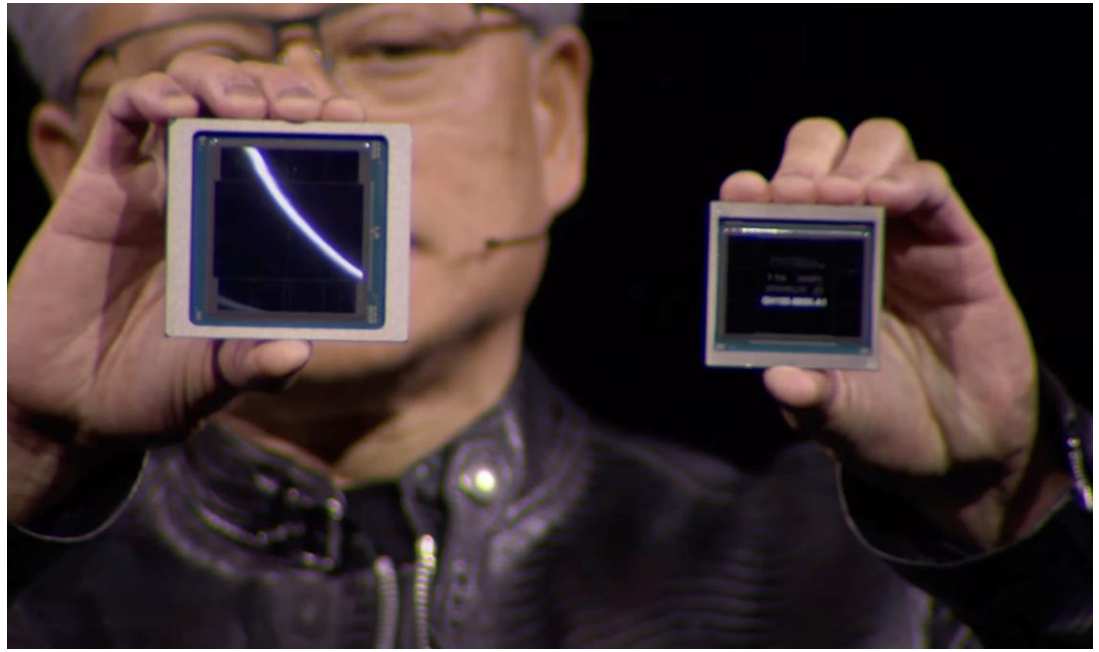

Nvidia Blackwell GB200 NVL72

- The GB200 NVL72, plugs 36 CPUs and 72 GPUs into a single liquid-cooled rack for a total of 720 petaflops of AI training performance or 1,440 petaflops (aka 1.4 exaflops) of inference. It has nearly two miles of cables inside, with 5,000 individual cables
- Amazon, Google, Microsoft, and Oracle are all already planning to offer the NVL72 racks in their cloud service offerings
- Nvidia says its systems can scale to tens of thousands of the GB200 superchips, connected together with 800Gbps networking with its new Quantum-X800 InfiniBand (for up to 144 connections) or Spectrum-X800 ethernet (for up to 64 connections)
- Nvidia has the DGX Superpod for DGX GB200, which combines eight systems in one for a total of 288 CPUs, 576 GPUs, 240TB of memory, and 11.5 exaflops of FP4 computing

Interesting AI Chip Design - IBM

- IBM NorthPole, building on TrueNorth Neuromorphic outperforms GPUs
- These chips represent a significant advancement in AI hardware, promising faster and more energy-efficient AI applications
- Brain-Inspired Efficiency: The NorthPole chip is 25 times more energy-efficient and has 22 times lower latency than comparable GPUs, thanks to its brain-like computation model
- Innovative Integration: TrueNorth integrates 4096 neurosynaptic cores, 1 million neurons, and 256 million synapses, offering a non-von Neumann, parallel, and low-power architecture
- Scalable Neural Networks: NorthPole supports larger neural networks by dividing them into smaller sub-networks, overcoming the limitation of not accessing external memory
- Event-Driven Architecture: TrueNorth has a Globally Asynchronous Locally Synchronous (GALS) design, enabling event-driven operations

Interesting AI Chip Design - Samsung

- **Mach-1 AI Accelerator Chip**:
  - Samsung is set to launch its own AI accelerator chip called **Mach-1** in **2025**
  - The goal is to break **Nvidia's** dominance in the AI accelerator market and reestablish Samsung as a leading semiconductor chip company
  - Mach-1 is specifically designed for **AI inferencing**
  - It aims to address the bottleneck in data transfer between the GPU, CPU, and memory by achieving data transfer speeds **1/8th** of current levels
  - Unlike some existing AI chips, Mach-1 can use **LPDDR memory** and is not solely dependent on High Bandwidth Memory
  - Samsung intends to compete with **AMD** and **Nvidia** using its in-house AI accelerator chip

- **Exynos Chips and AI Design (Europe?)**:
  - Samsung has been using **artificial intelligence (AI)** to design its **Exynos** chips
  - The company collaborates with **Synopsys**, a leading chip design software firm, to create its new chips
  - By leveraging AI and advanced software, Samsung aims to enhance the performance and efficiency of its Exynos processors

- **HyperCLOVA X AI Chip**:
  - Samsung and Naver jointly developed an AI chip that is **8 times more power-efficient** than Nvidia's
  - This chip, revealed as a **Field-Programmable Gate Array (FPGA)**, is used to power Naver's mega-scale AI model called **HyperCLOVA X**
  - Developers can make changes to the chip's design during the prototype phase before mass production

Quantum Chips

- Many competitors
- Trying to outperform classical computers
- Some tasks can’t be done on a classical computer: materials, weather, IoT

- IBM Heron, Osprey(433) & others
- Intel Tunnel Falls, Tangle Lake(49)
- Google Sycamore, Bristlecone(72)
- Microsoft Azure Quantum
- China Jiuzhang(76)
- Honeywell

A Few Sample Products

COMPUTERS

HP Dragonfly G4 Specifications

- **CPU - 13th Gen Intel Core processor**
- **GPU - Intel integrated**
- **RAM - Up to 32 GB**
- **Storage - Up to 2 TB**
- **Display - 13.5-inch, 3:2, 3K2K OLED, 400 Nits**
- **Camera - 5MP**
- **Audio - Bang & Olufsen, dual mics**
- **Ports - 2x Thunderbolt 4 with USB C, 2x SuperSpeed USB Type-A, HDMI 2.0, Nano-SIM, Audio jack**
- **Connectivity - Wi-Fi 6E, Bluetooth 5.3, 4G LTE or 5G (optional)**
- **Dimensions - 12.39 in x 8.78 in x 0.72 in**
- **Weight - Starting at 2.98 lbs.**
- **$2,031**

Dell

XPS 13-16

- 14th Gen Intel Core Ultra
- Dedicated GPU (NVIDIA/ARC)
- 500 Nits
- $1,300-3,350

Latitude 13-16

- 14th Gen Intel Core Ultra
- Integrated GPU or ARC
- 250 Nits
- $1,289-2,369

Lenovo

X1 2-in-1

- Think 2-in-1 as a form factor = laptop and tablet in one device
- Laptop, tablet, & tent modes
- $3,759

X1 Carbon

- Continues to be leading edge
- Intel® Core™ Ultra 7 in Gen 12
- Gen 4 SSD
- 400 nits
- $1,803-3,069

Apple MacBook Pro

- **CPU – M3 Pro 11, 12 or Max, 12, 14, 16 Core – March 8**
- **GPU – Pro (14), 18, 30, or 40 Cores, 16-core Neural Engine**
- **RAM - Up to 96 GB**
- **Magic Keyboard with Touch ID, Force Touch trackpad**
- **Storage - Up to 8 TB**
- **Display –14, 16” Liquid Retina XDR display², 1600 nits**
- **Camera - 1080p FaceTime HD camera**
- **Audio - High-fidelity six-speaker sound system with force-cancelling woofers**
- **Ports - Three Thunderbolt 4 ports, HDMI port, SDXC card slot, headphone jack, MagSafe 3 port**
- **Connectivity - Wi-Fi 6E, Bluetooth 5.3**
- **Dimensions - 12.31 in x 8.71 in x 0.61 in**
- **Weight - Starting at 3.6 lbs.**
- **$1,599-4,199**

Surface Pro Copilot Plus PCs

- 11th Edition with NPU, with Recall intelligence for local AI
- Snapdragon X Elite Processor
- PRISM enables apps built for an Intel processor to run
- OLED in higher end models
- Haptic Flex keyboard
- 14 to 22-hour battery life
- 2 USB-C, Wi-Fi 7, < 2 lbs.
- $1,000-1,999 on June 18

Framework Laptop Specifications

- **CPU – Ryzen 7 or 9, Intel 13th Gen**
- **GPU – Radeon 780M or RX7700S**
- **RAM - Up to 64 GB**
- **10 key can be on the right or left**
- **Storage - Up to 4 TB**
- **Display – 13, 16”, 500 nits**
- **6x user-selectable expansion cards**
- **Connectivity - Wi-Fi 6E, Bluetooth 5.2**
- **Dimensions – 14.0 in x 11.4 in x 0.70 in**
- **Weight – 2.1 – 2.4 kg**
- **$849-2,053**

Ricoh Scanners

1. **Ricoh fi-8170**:
   - A **high-speed desktop scanner** designed for reliable color scanning
   - Estimated daily scan volume of **10,000 sheets**
   - Features **Automatic Separation Control** and **Multi-feed Detection** to prevent jams
   - Ideal for **business environments** with its robust performance and compact footprint
2. **Ricoh fi-8040**:
   - A **front-office scanner** with network scan capabilities
   - **PC-less scanning** directly to email, network folders, or FTP servers
   - **Wi-Fi-enabled** for flexible placement
   - **40 ppm** scanning speed and **600 DPI** optical resolution
3. **ScanSnap iX1600**:
   - A **Wi-Fi-enabled color network scanner**
   - **One-touch scan and store** to up to 30 destinations, including cloud applications
   - **60 ppm** scanning speed and **600 DPI** optical resolution
   - Comes with **Adobe Acrobat Pro DC** subscription
4. **Ricoh fi-7300nx**:
   - A **fast Wi-Fi-enabled color network scanner**
   - **Ideal for high-traffic areas** with its compact footprint
   - **60 ppm** scanning speed and **600 DPI** optical resolution
   - **Cloud-ready** and supports **TWAIN and ISIS drivers**

Ricoh Enterprise Scanners

fi-8950 and fi-8930

$14,500 & $9,900

Powerful, Feature-Rich Solutions For Exceptional Long-term Value

fi-8950	fi-8930
150 ppm/300 ipm	130 ppm/260 ipm
750 sheets (A4)/500 sheets (A3)	750 sheets (A4)/500 sheets (A3)
Daily volume: 130,000 sheets	Daily volume: 110,000 sheets
Automatic skew correction	Automatic skew correction
Stapled document detection	Stapled document detection
Clear image capture	Clear image capture
Paper ejection control	Paper ejection control
Touchscreen and button interfaces	Touchscreen and button interfaces
Straight scan capabilities	Straight scan capabilities
Lifetime of consumables: 700,000 sheets	Lifetime of consumables: 700,000 sheets

Canon Scanners

1. **Canon image FORMULA R40**: This tabletop scanner offers excellent image quality, a small footprint, and a user-friendly app. It has a **50-page automatic document feeder** and is priced around **$250**.
2. **Canon CanoScan LiDE 300**: A budget-friendly option with a slim design. It's easy to use and provides decent scanning performance at a price of approximately **$58**.
3. **Canon DR-C230**: An entry-level scanner suitable for basic scanning needs
4. **Canon DR-C225**: Known for its speed and auto-feed capabilities
5. **Canon DR-M160II**: A professional-grade scanner with high-quality results
6. **Canon DR-M260**: Designed for high-volume scanning tasks
7. **Canon CR-L1**: A versatile check scanner for Remote Deposit Capture (RDC)

Apple iPhone 16, Plus, Pro, & Max

- All models on the A18 Pro or Bionic instead of the A17
- 3nm – N3E from TSMC
- Neural Engine core increase
- Geekbench 6 Speeds - Single/multi core
  - A18 – 3500/8200
  - A17 Pro – 2906/7231
  - Snapdragon 8 Gen 4 – 2985/10762

Apple iPhone 16, Plus, Pro, & Max

Features	iPhone 16	iPhone 16 Plus	iPhone 16 Pro	iPhone 16 Pro Max
Price	$799	$899	$999	$1,199
Display – Super Retina	6.1-in XDR OLED	6.7-in XDR OLED	6.3-in XDR OLED	6.9-in XDR OLED
Processor	A18	A18	Apple A18 Pro, 1TB	Apple A18 Pro, 1TB
Resolution	2446 x 117 pixel	2796x 2290 pixel	2622 x 1290 pixel	2868 x 1320 pixel
Camera 48 MP Fusion	12MP UW	12MP UW	Wide 12MP Tele 5x | 48MP UW | 12MP Selfie	Wide 12MP Tele 5x | 48MP UW | 12MP Selfie

Apple Watch Series 10 & Ultra 2

- The Apple Watch Series 10 packs a swifter S10 chipset than the Series 9's S9 silicon
- Always on display at 2,000 or 3,000 nits
- Sleep Apnea feature
- $249-799

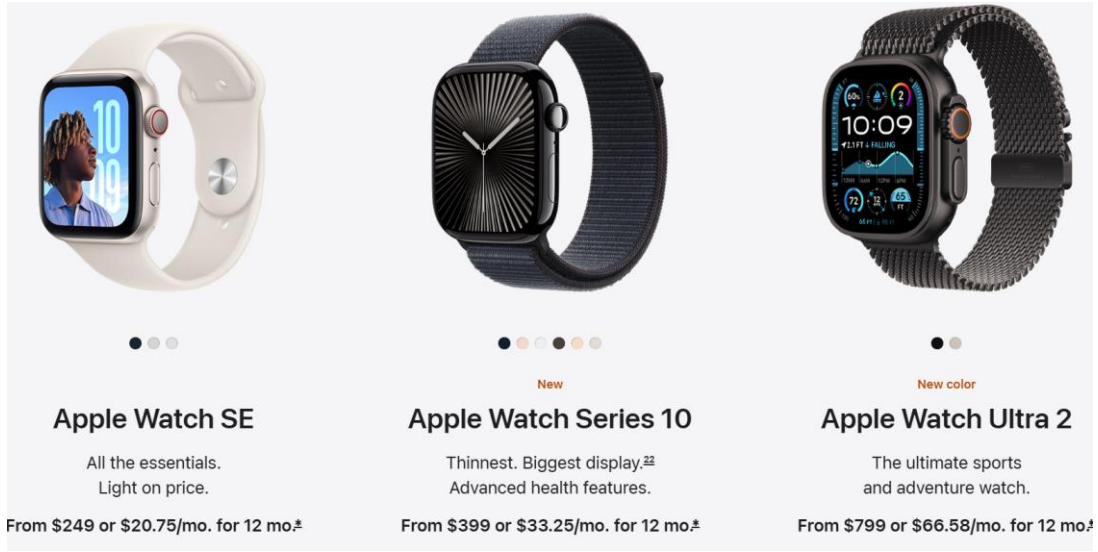

Apple AirPods 4 Are The Big News

- Up to 2x more Active Noise Cancellation, with Adaptive Audio, and Transparency mode
- Hearing Test, Hearing Aid, and Hearing Protection features (coming fall 2024)
- Conversation awareness
- H2 chip
- $129, 179, 249, and 549

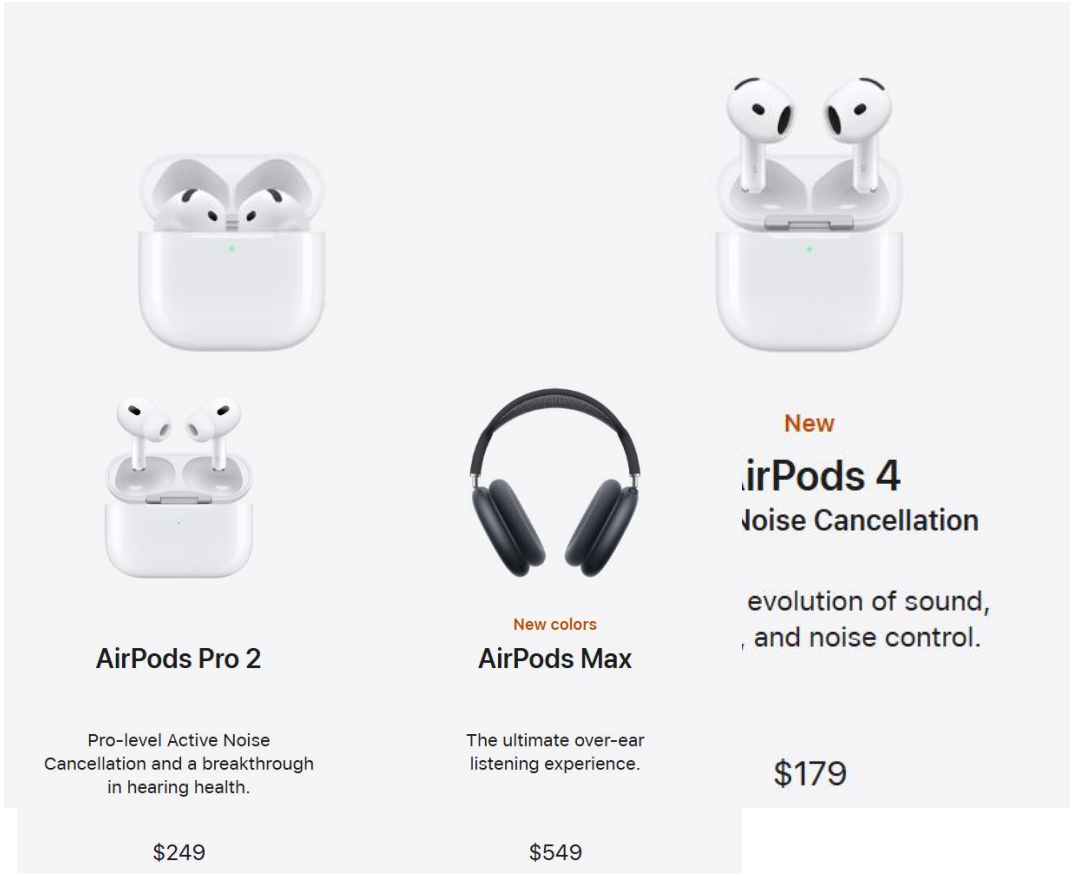

Finally, Some New iPads!

- 13” iPad Air - $799+
- iPad Pro gets
  - An OLED display
  - M4 processor – 3nm 10 core
  - 11” $999, 13” $1,299+
- Apple Pencil Pro is haptic
- New Magic Keyboard $349

Apple Home Robotics

- Large iPad-like display mounted on a "thin robotic arm“
- Tilt and up and down and rotate a full 360 degrees
- Smart home command center
- Videoconferencing machine such as for FaceTime calls
- Home security monitoring tool
- 2026 or 2027
- Around $1,000

Apple Vision Pro “Spatial Computing”

- **M2 and brand new R1 chips for the sensor inputs, 12ms**
- **Twin 4K displays inside the headset for HDR picture quality. For glasses you'll need custom $100 ZEISS inserts**
- **EyeSight uses a front display to show others that you are engaging with them or are immersed in the headset**
- **visionOS allows for 3D experiences native to the Vision Pro. iOS and iPadOS apps are available in a Vision Pro App Store along with visionOS apps**
- **The crown on the top of the headset allows you to control how immersed in the headset you are, from augmented reality to virtual reality.**
- **No controllers needed — the headset relies on eye and hand tracking in addition to voice commands**
- **$3,499-5,000, costing $1,509 to build**
- **January 19, 2024**

Google Pixel 9/Pro/Pro XL/Pro Fold

- Google Tensor G4 Titan M2 security coprocessor “Zuma Pro”
- Likely G5 will make it for Pixel 10 next year
- Samsung Modem 5400 for 5G non-terrestrial networks (NTN)
- Triple camera setup
- Four models
- August 14, 2024
- $799-1,799

More From The Pixel Portfolio

New

Pixel Buds Pro 2

Light ears ahead.

$229

or $19.08/month with 12-month financing*

New

Google Pixel Watch 3

Guidance. Goals. Gorgeous.

From $349.99

or $29.17/month with 12-month financing*

Pixel Tablet

Help in your hand. And at home.

From $399

or $33.25/month with 12-month financing*

Samsung Galaxy Z Fold 6 & Z Flip 6

- **Announced July 10 to promote during the Olympics, negotiated 18 months in advance**
- **STF coating on screen with dual rail hinge**
- **Fold Colors: Pink, Silver Shadow, Navy; Exclusive colors on Samsung's website: Crafted Black, White**
- **Flip Colors: Silver Shadow, Yellow, Blue, Mint; Exclusive colors on Samsung's website: Crafted Black, White, and Peach**
- **Snapdragon 8 Gen 3, 12 GB RAM**
- **Wider cover screen on both**
- **S Pen included**
- **Dust resistance**
- **No more crease**
- **July 24, 2024**
- **$1,100-1,900**

Samsung Watch Ultra / 7

- Three-nanometer application processor
- 100-hour battery life
- Dual-frequency GPS system
- Also Watch 7 with biometrics and health for business professionals

Samsung Galaxy Ring

- Health App
- Sleep Activity
- Heart Rate
- Titanium
- Seven days on a single charge
- No subscription required

Galaxy S24/S24+/S24 Ultra

- Snapdragon 8 Gen 3
- AI Translation in 12 languages
- 50MP Rear Zoom, 10-200 MP wide angle, 12MP ultrawide angle, and 12 MP Forward-Facing Camera
- The phones may also feature a 10MP short lens
- 128 GB, 256Gb, 512GB and 1TB storage
- 6.8” QHD+ or 6.2” FHD+ AMOLED displays, higher refresh rate
- 5,000mAh energy cell
- Four colors: Titanium Gray, Black, Violet, and Yellow (replacing green)
- January 17, 2024

Galaxy S25/S25+/S25 Ultra

- Camera Improvements
- Entry level size to 6.36” from 6.2 on the S24
- Ultra screen size 6.9” from 6.8”
- Screen brightness increase
- Wide angle lens upgrade to 12 or 50MP
- Variable telephoto
- Snapdragon 8 Gen 4, matching the A18 numbers, or X Elite?
- Android 15

Humane Ai Pin

- $699, $24/mo. subscription

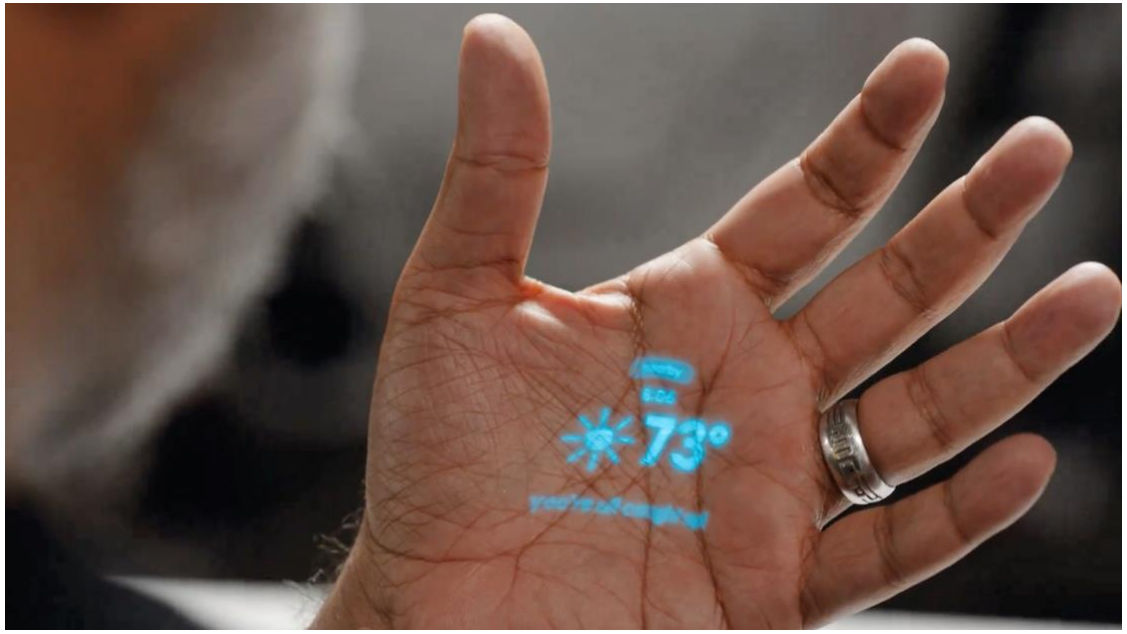

- Post-smartphone device

Important Technologies & Standards To Follow

COMMUNICATIONS & STANDARDS

Wi-Fi 7 Unifi U7 Pro & Three More

- The U7 Pro is a ceiling-mounted Wi-Fi 7 access point designed for demanding, large-scale environments
- It supports 6 spatial streams and operates in the 6 GHz frequency band for interference-free Wi-Fi
- With a coverage area of 140 m² (1,500 ft²), it can handle 300+ connected devices
- Powered via PoE+ and equipped with a 2.5 GbE uplink, it’s suitable for high-performance networks
- Plus U7 Pro Max, Pro Wall, Outdoor
- $189, $279, $199, $199

Wi-Fi 7: What's New

- Newest Wi-Fi standard now out (802.11be)
- Speeds up to 46 Gbps
  - Unless using for mesh, need ethernet speeds above 1 Gbps
- Uses same frequencies as Wi-Fi 6E, but allows for channels of up to 320 MHz vs. 160 MHz for 6/6E
- All three standards use WPA 3, which encrypts all data between the access point and devices
- The first devices are available from Unifi and TP-Link now

	Wi-Fi 6	Wi-Fi 6E	Wi-Fi 7
IEEE Standard	802.11ax	802.11ax	**802.11be**
Frequencies used (GHz)	2.4, 5	2.4, 5, **6**	2.4, 5, 6
MIMO Streams	8	8	**16**
Channel Width	20, 40, 80, 80+80, 160MHz	20, 40 , 80, 80+80, 160MHz	20, 40 , 80, 80+80, 160, 320MHz
Max speed	9.6 Gbps	9.6 Gbps	**46 Gbps**
WPA Standard	WPA 3	WPA 3	WPA 3

Wi-Fi 7: Technical Considerations

- Wi-Fi 7 includes 320 MHz Channels & 4K QAM (**Quadrature Amplitude Modulation**), providing more densely packed channels and more data density.
- The benefit of a typical Wi-Fi 7 laptop is a potential maximum data rate of almost 5.8 Gbps. This is 2.4X faster than the 2.4 Gbps possible with Wi-Fi 6/6E and could easily enable high-quality 8K video streaming. In addition, it could reduce the time for a 15 GB file download to roughly 25 seconds vs. the one minute it would take with the best legacy Wi-Fi technology.
- With MLO (**Multi-Link Operation**), Wi-Fi 7 devices can simultaneously connect on two bands. This enables faster speeds through aggregation. Or, both bands can be used concurrently to share redundant/unique data for improved reliability with ultra-low and precise latencies.

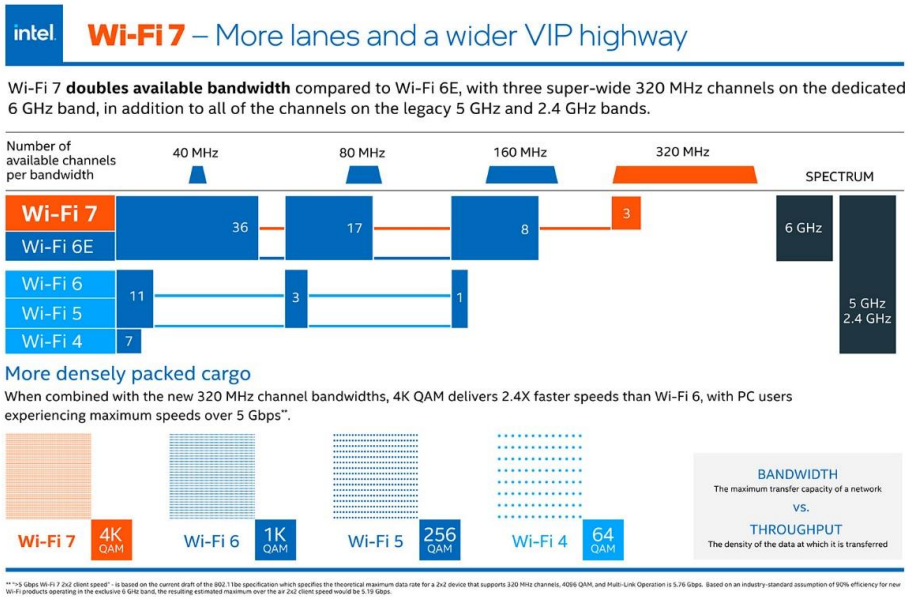

New LEO For Phones

- AT&T, T-Mobile, Google, Samsung, Apple all participate
- Starlink sent and received texts over a 4G/LTE connection between mobile phones via its latest generation of satellites, called v2mini, for the first time in January 2024, following similar projects from Amazon, Apple, AST SpaceMobile, Huawei, and Lynk Global
- Factors: Larger dishes, beamforming, satellites lower

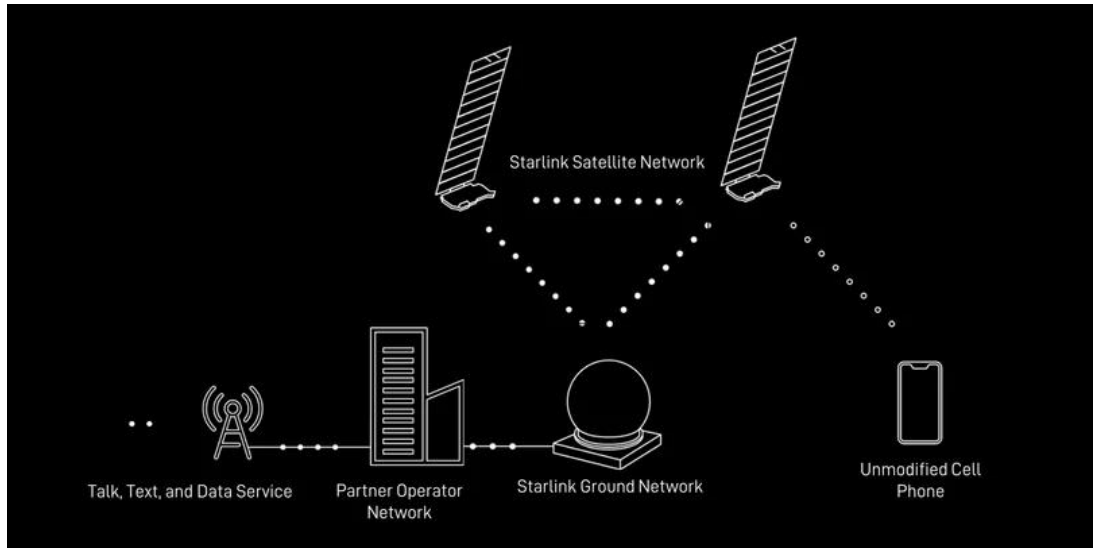

6G And Beyond

- Extreme High Frequency (EHF)
- 6G will deliver theoretical terabit speeds
- Most users will get more than 100Gbps
- Extreme-low latency of less than one millisecond
- 6G satellite technology and intelligent surfaces capable of reflecting electromagnetic signals will deliver low latency, multi-gigabit connectivity to parts of the world where it has been too difficult or too expensive to reach
- Early 2028, but could be as late as 2030

Faster Ethernet

- **100 Gbps** Ethernet is now commonplace in data centers, enabling rapid data transfer
- **IEEE 802.3bz** (2.5 Gbps and 5 Gbps) provides intermediate speeds for enterprise networks
- **IEEE 802.3cd** (50 Gbps and 200 Gbps) targets high-speed data center links
- **IEEE 802.3cm** (400 Gbps) is under development
- Challenges and Research Areas: Signal Integrity, Energy Efficiency, Optical Ethernet, Terahertz Ethernet

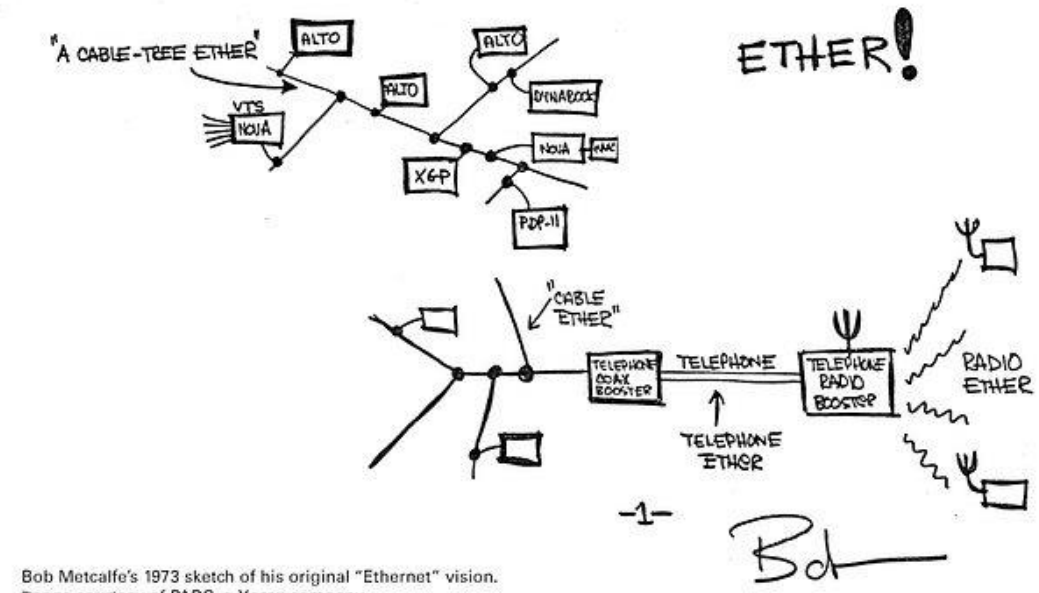

Bob Metcalfe's 1973 sketch of his original "Ethernet" vision.
Image courtesy of PARC, a Xerox company

FCC Quadruples Basic Broadband

- Standard was updated March 14, 2024, last updated 2015
- 100 Mbps download and 20 Mbps upload is the new standard
- Four-fold increase compared to a 2015, which set the benchmark to 25 Mbps download and 3 Mbps upload
- Goal is to have 1 Gbps/500 Mbps for American consumers
- Benchmark compliance is measured with NTIA's BEAD Program and multiple FCC USF programs
- Nvidia GeForce Now game streaming service requires at least 35Mbps for 1080p at 240 FPS
- For movie streaming, Netflix needs 15 Mbps or higher for 4K streaming but it also aspires to expand into cloud gaming

Other Evaluations By The FCC

- Last year, the White House announced plans to invest $42 Billion for universal internet access by 2030
- Fixed terrestrial broadband service (excluding satellite) has not been physically deployed to approximately 24 million Americans, including almost 28% of Americans in rural areas, and more than 23% of people living on Tribal lands
- Mobile 5G-NR coverage has not been physically deployed at minimum speeds of 35/3 Mbps to roughly 9% of all Americans, to almost 36% of Americans in rural areas, and to more than 20% of people living on Tribal lands
- 45 million Americans lack access to both 100/20 Mbps fixed service and 35/3 Mbps mobile 5G-NR service
- Based on the new 1 Gbps per 1,000 students and staff short-term benchmark for schools and classrooms, 74% of school districts meet this goal
- For casual users, 100 Mbps should be plenty
- Net Neutrality reviewed on April 25

Linux Back?

- If everything is Cloud/SaaS, can we lower cost/risk locally?
- Parallels DaaS
- Inuvika
- CrossOver – Codeweavers
- The LAMP Stack (Linux, Apache, MySQL, PHP)is a classic open-source stack that's versatile and cost-effective and ideal for basic web applications
- Licensing costs frequently approach zero with improved security

Besides Becoming The Largest Company, Much More Is Expected

MICROSOFT OBSERVATIONS

Microsoft Has Passed Apple As The World's Most Valuable Firm

Three Divisions For Growth

- **Azure -** closing on Amazon, 25% of $212B revenue, grew 27% YOY, gross margins of 60%
- **Microsoft 365** - $20B, 10%/yr.
- **Video Games** - $15B will grow with Activision Blizzard acquisition

AI Advantage

- Since 2014, margins have risen from 28% to 43%
- Copilot 365 assembling email and spreadsheets is much easier than other AI tasks such as Copilot on GitHub and has a 52-83% markup
- First mover
- Building their own AI chips for Azure
- Had the OpenAI team as employees, and then they didn't

Windows 12

- **Next OS Iteration:** Windows 12 is the next major version of the Windows operating system, not likely in 2024
- **Redesigned Interface:** Anticipate visual enhancements, streamlined menus, and a more intuitive user experience
- **Performance Boost:** Improved performance, faster boot times, and optimized resource utilization
- **Enhanced Security:** Robust security features, including built-in protections against threats
- **Compatibility:** Focus on maintaining compatibility with existing software and hardware
- Free upgrade, but outright purchase price between $139-169

Office 2024

- **Perpetual License:** Office 2024 is a ‘perpetual license’ version, meaning a single payment for ongoing use without annual subscriptions. It replaces Office 2021
- **Office LTSC 2024:** Similar software for volume license customers, following the Long-Term Servicing Channel (LTSC) model
- **Enhanced Features:** Expect improvements in productivity tools, collaboration features, and stability
- **Platform Agnostic:** Available for both Windows and Mac users
- **Continued Support:** Office 2024 will receive long-term support and updates
- Office 2021 was $439.99, Office 2024 will likely be $499
- If you have Microsoft 365, DON’T buy Office 2019, 2021, or 2024

Microsoft’s Power Platform

The low-code platform that spans Office 365, Azure, Dynamics 365, and standalone applications

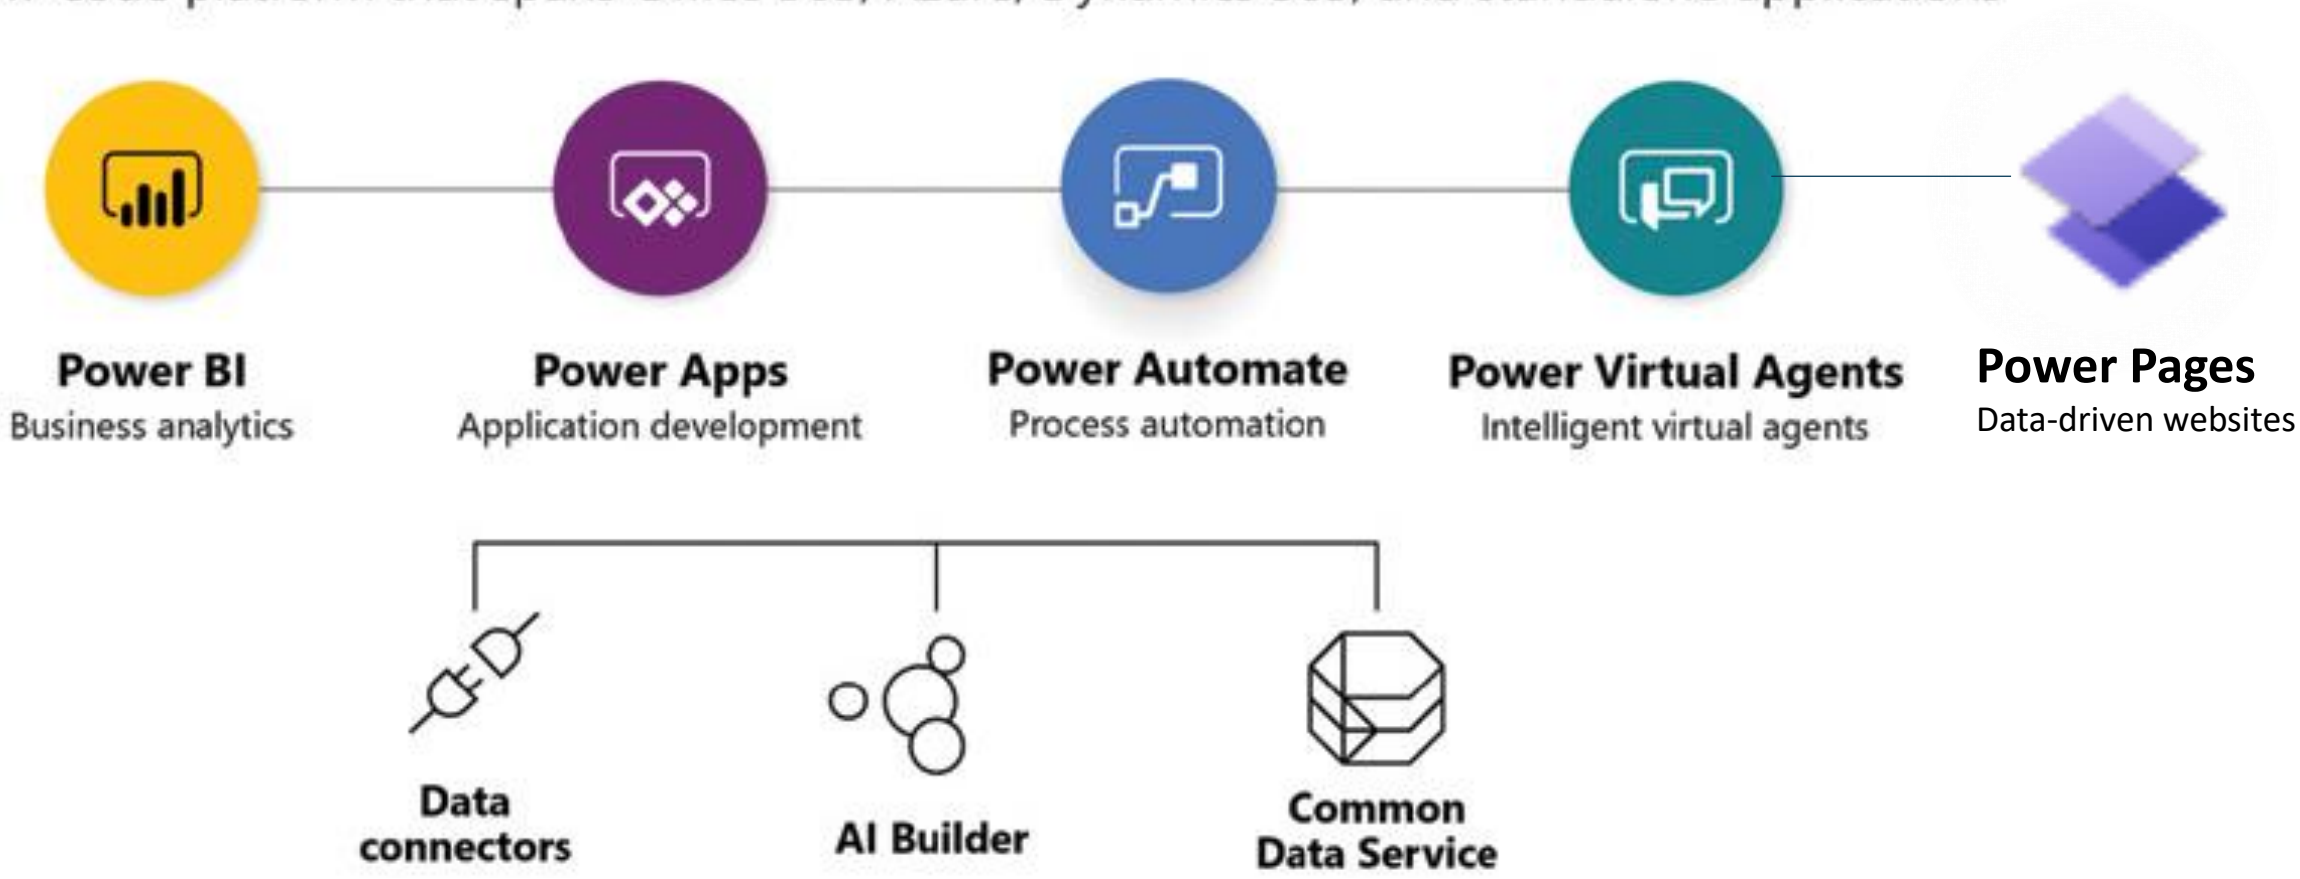

Source: Microsoft

Microsoft Copilot

- **AI Companion:** Copilot is your everyday AI companion, designed to bring the power of generative AI to everyone across work and life
- **Expanded Availability:** It’s now available across devices, through the web, or via the Copilot mobile app on iOS and Android
- **Free Web Apps:** Copilot is part of your Copilot Pro subscription in the free Microsoft 365 web apps, including Word, Excel, PowerPoint, OneNote, and Outlook
- **Custom GPTs:** You can create personalized Copilot GPTs to assist with specific tasks based on your interests, like a career counselor or study buddy
- **Global Reach:** Copilot Pro is available in all 222 countries/regions where Copilot is accessible, empowering power users and creators worldwide

Microsoft’s AI Copilots

Name	Products	Monthly Cost	Commercial Data Protection Included?
Copilot for MS 365 (Business/Enterprise)	Microsoft 365 apps (Word, Excel, PowerPoint, Outlook, Teams)	$30 per user	Yes
Copilot in Windows (Bing Chat)	Windows OS	Free	Not for home users, included with most business/enterprise O365/M365 plans
Copilot Pro for Individuals	Advanced features on top of standard Copilot, plus integration with home Microsoft 365 apps	$20/user/mo.	Not specified
Copilot for Security	Microsoft’s cybersecurity products	Consumption-based fee - $4/hour	Not specified
Copilot for Finance, Sales, and Service	Financial operations, sales optimizations, service enhancements	$50/user/mo., $20/user/mo. if already have MS 365	Not available to other customers, runs on Microsoft cloud in separate instance of ChatGPT, not used by MS to train models by default
Designer for Copilot	Image creation and editing	Not available	No
Copilot GPTs and Azure AI Studio	Custom generative AI assistants and solutions	Not specified	Not specified

Microsoft Security Copilot

- Available April 1, 2024
- $4/hour for SCUs (Security Compute Units)
- Conversational skills that can write incident reports quickly
- AI assistant for
  - Security operations tasks
  - Identity-related tasks
- Governance for AI (NIST)

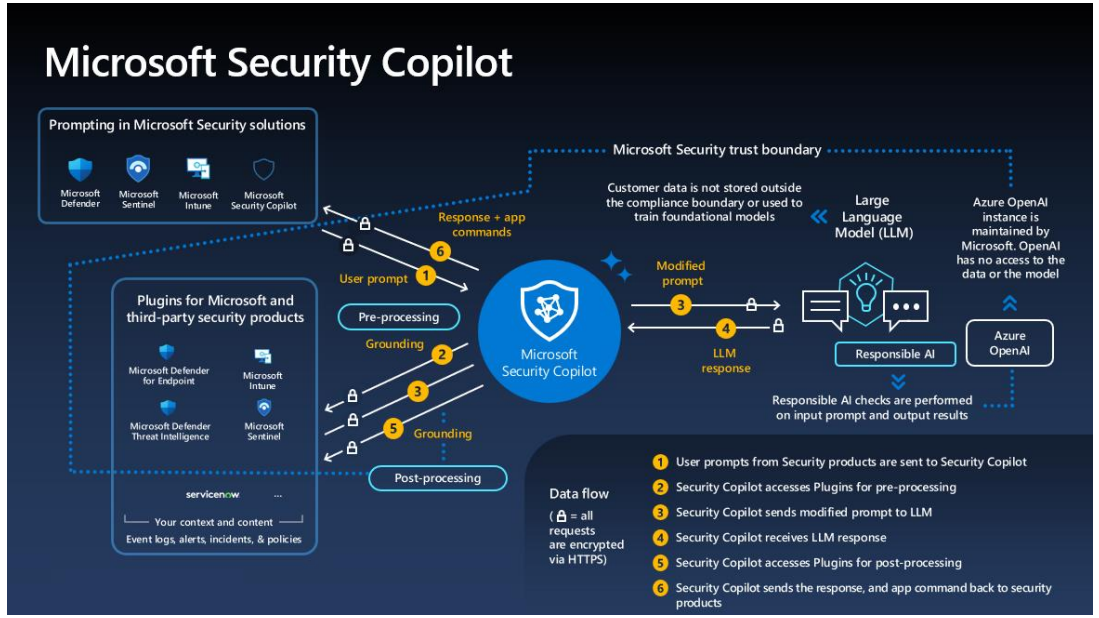

Private Equity, Supply Chain, and Global Elections

EXTERNAL MARKET FACTORS

PE Affects On Profession/Industry

Business Consequences

- Lack of investment/vision in many platforms
- Marketing numbers often don’t reflect reality
- Support frequently reduced
- Rarely add much value
- Much “dry powder” capital
- I’m told to “get over it”

Personal Impacts

- Good exits for partners and owners
- Expectations of team members and employees increased
- Five-year cycles with bonuses
- Numbers managed “from a spreadsheet”
- Transactional

Supply Chains Remain Tight

- Demand likely to exceed production
- Planners conservative with inflation uncertainty, wars & election year, where 4 Billion people in 50 countries will have elections
- Manufacturing plants around the world not coming online as quickly as expected
- PC demand up

Regulatory, Tools, And Policies

MANAGEMENT AND GOVERNANCE

8 IT Failures Of 2023 - CIO

- Tech trouble in the skies
  - United, Hawaiian, and Southwest had service outages
  - But so did the FAA
- The NYSE's brittle backup process
  - Incorrect backup to Chicago
  - Caused the next day session to continue the previous day's session
- In space, no-one can cancel your software license
  - Oracle software use for Space Shuttle has $15M price over 3 years
  - Can't shut it off, can't afford audit
- Software licensing situation cloudy
  - Nutanix used software it didn't license
- Turn off the lights, the party's over
  - Minnechaug Regional High School in Massachusetts had been happily running a "green lighting" system installed by 5th Light, hit by malware, stayed on at full brightness continuously for ten years
- When a crash means a real crash
  - MRH-90 Taipan, a military helicopter used in Australia, not patching led to training crash
- Cascading phone failures down under
  - Optus outage from routing changes by Singtel
- Artificial intelligence, real failure
  - Lawyers at Levidow, Levidow & Oberman using ChatGPT with court case hallucinations
  - CNET retracted 35 stories written with RAMP (Responsible AI Machine Partner)

Governance

Policies

- Responsible AI Institute
- FinCEN Guidance (BOI)
- Integrated Growth Advisors
- Supervizor – Continuous Audit
- TimeCredit – Technical Memos

Privacy

- Privacy Lock – Global Privacy Compliance
- NIST updated its Cybersecurity Framework on February 26, 2024, for the first time in a decade, adding a new area of focus: governance

SaaS Management

Costs

- Service scalable on demand
- Variable performance for users
- SaaS licensing tools
  - Privacy.com (Virtual Card for SaaS payment restrictions)
  - Josys
  - Productiv
  - StackShine (now owned by FinQuery(formerly LeaseQuery))
  - Spendflo
  - Substly (SMB)
  - Zylo

Security

- Distributed Endpoints
- Security risk increase
  - EDR/MDR
  - Unified Endpoint Management
- Security risk decrease
  - Cloud usually more secure
  - Silos isolate risk by application

Effective Backup Plan

- Data loss can cripple an organization or put it out of business altogether
- An effective backup plan should be:
  - Automatic
  - Redundant
  - Regularly tested
  - With a copy stored off-site
  - With a local isolated or air-gapped copy
- 3-2-1 – 3 backup copies, 2 media types, 1 copy offsite

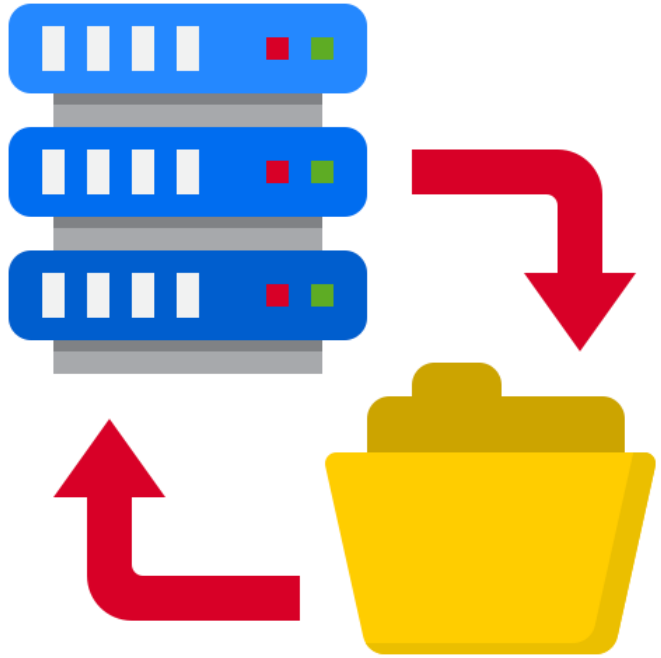

Isolated And Offline Backup

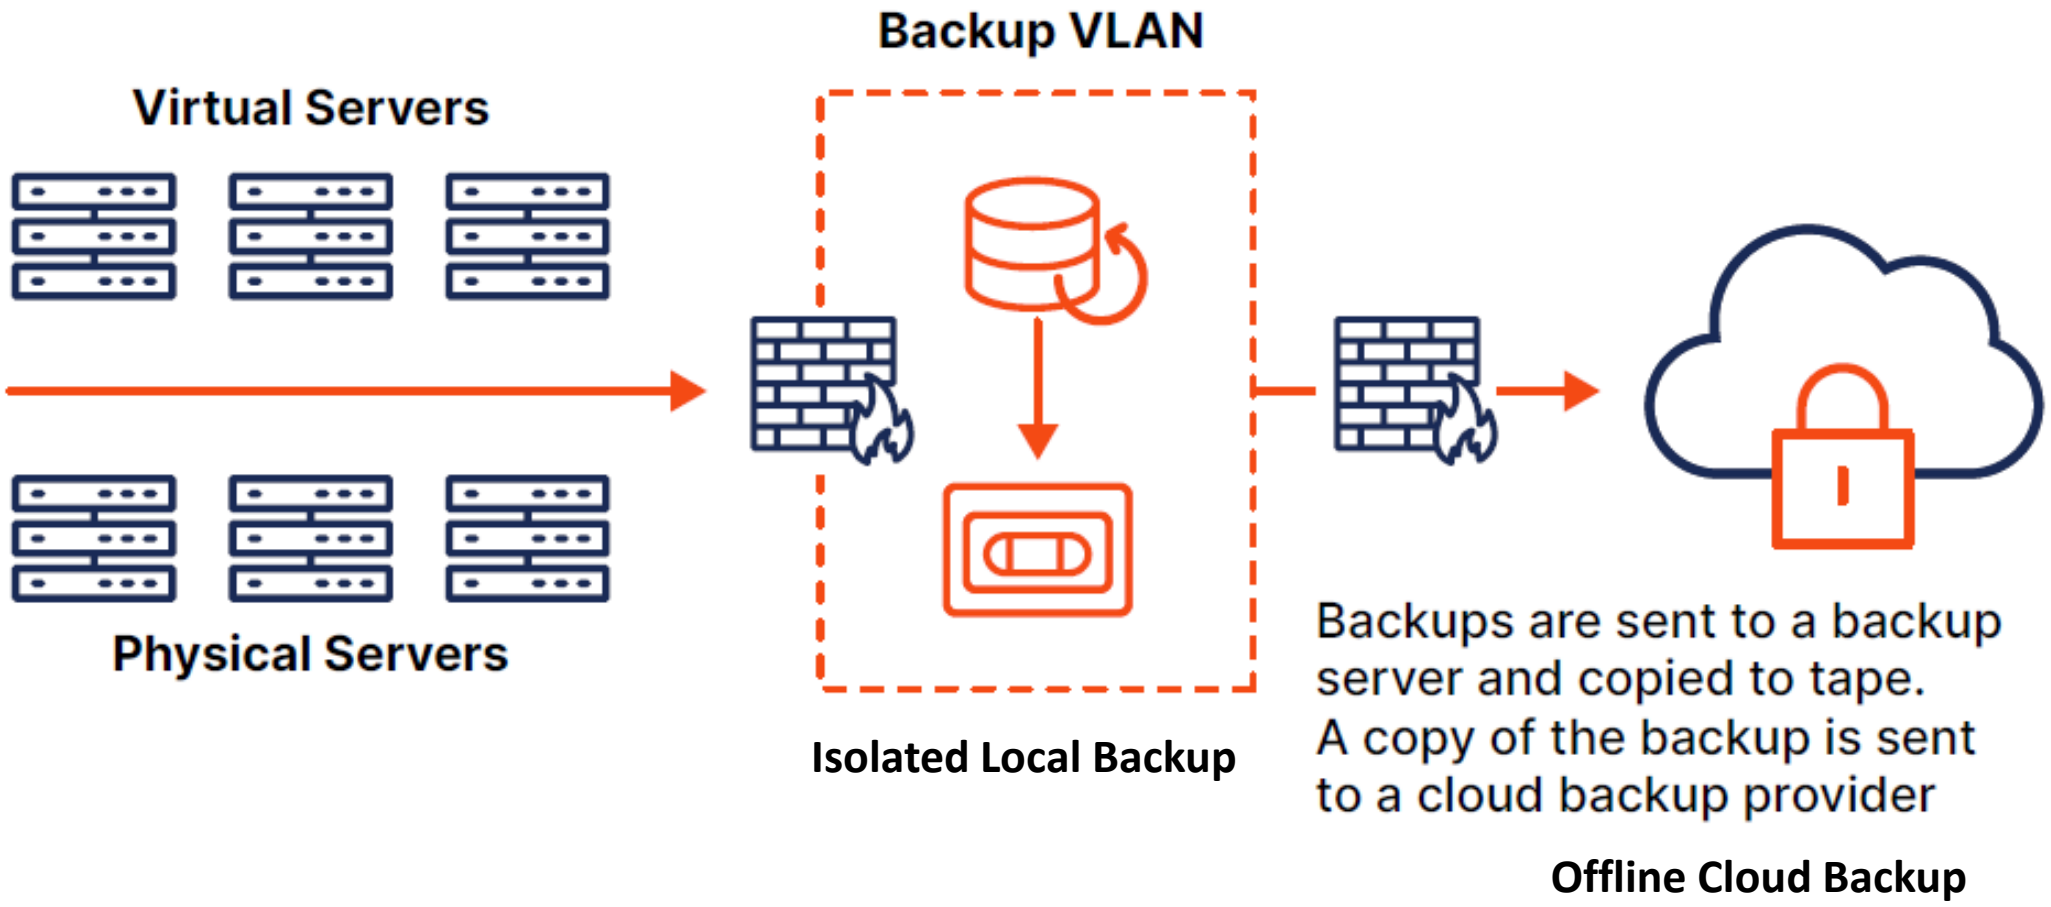

Source: Liska, *Ransomware*, ActualTech

Top Ransomware Groups (July 22 - June 23)

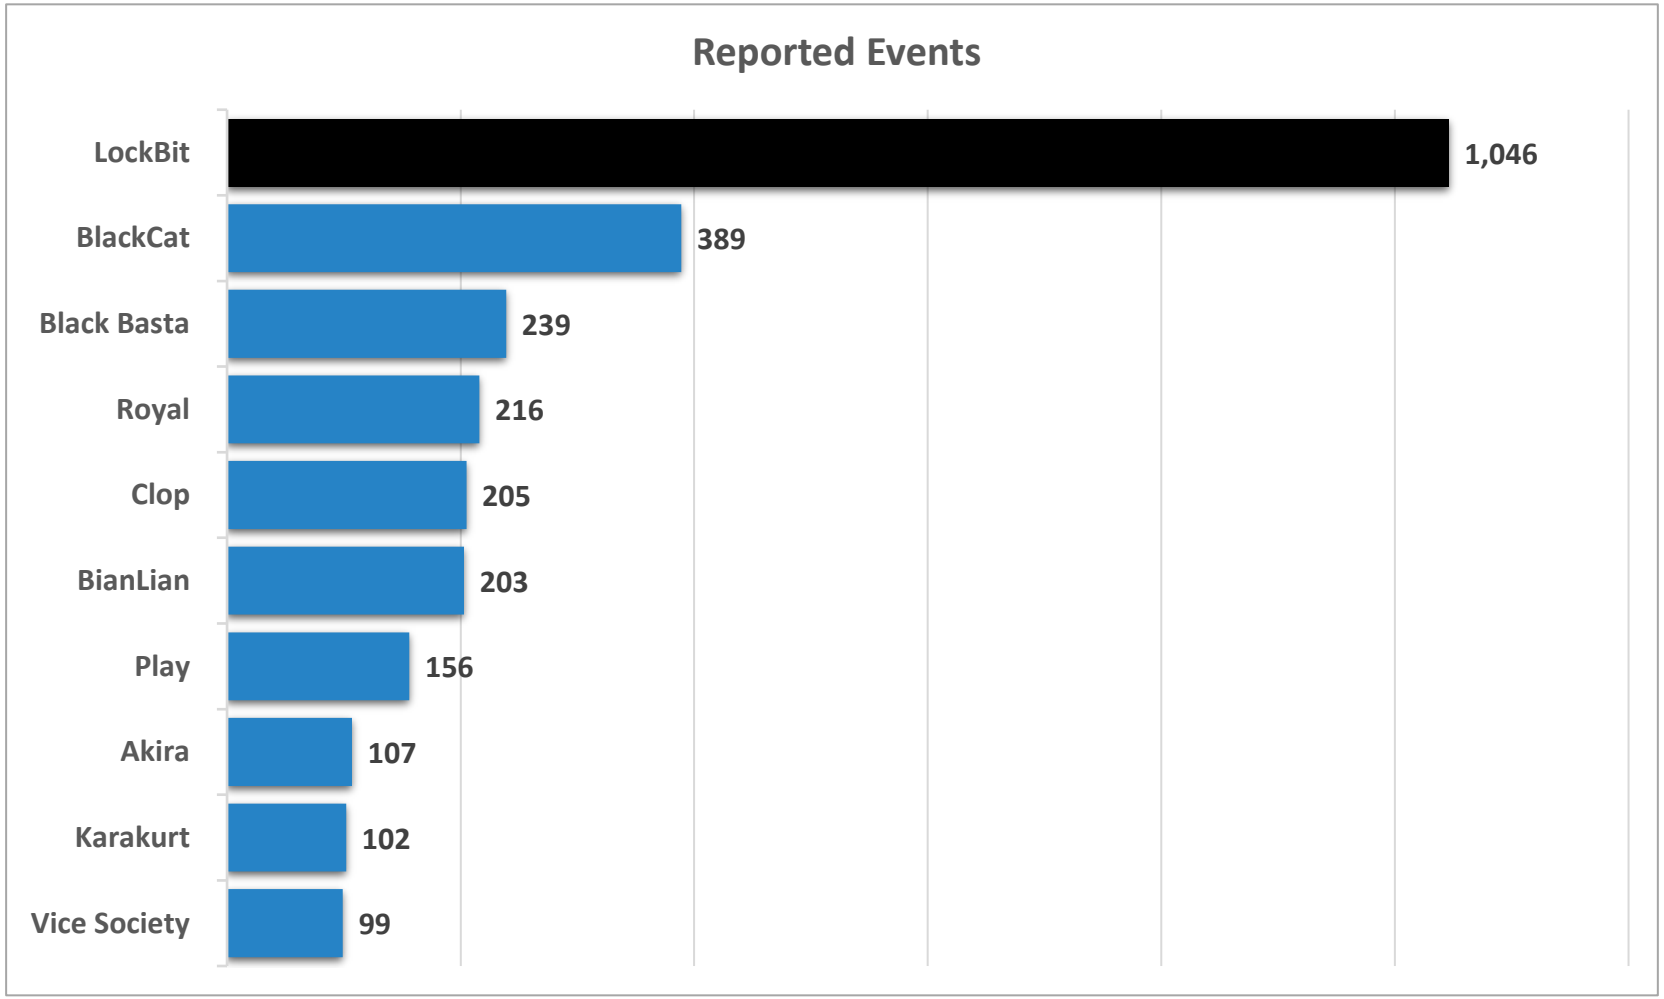

Source: Tom's Hardware

Ransomware Cryptocurrency Payments

Bitcoin accounts for 98% of ransomware payments

Cumulative Ransomware Payments, 2023 vs. 2022
Through June 2023

Evolving Cybercrime Interdependence

- Cybercrime is evolving into a complex web of interdependent cybercriminals
- Cybercriminals are selling their knowledge as a service to other cybercriminals
- Types of services offered
  - Internet access brokers (IABs) – acquire and sell compromised credentials
  - Counter anti-virus services (CAVs) – scan malware to identify flagged code before initiating a malware campaign
  - Negotiation customer services – negotiate payments with victims
  - Bulletproof hosting – information of cybercriminal customers is not collected or disseminated

2023 Europol IOCTA Report
Internet Organized Crime Threat Assessment

Social Engineering

- A common form of social engineering is "vishing" (voice phishing)
  - For example, a person pretending to be an IT manager may request your username and password to perform required server maintenance
  - Or a person, using information gleaned from LinkedIn or Facebook, may impersonate a user to Help Desk personnel to recover or reset a password
- More than *90 percent* of all cyberattacks start with phishing
  - Targeted phishing attacks that include phone calls are three times more effective, as reported by IBM in 2022

Deepfake Audio

- UK CEO received a spear-phishing e-mail from someone posing as his parent company's CEO
- Message requested that $302,000 be wired in the next hour to a Hungarian supplier
- The e-mail was followed up with a phone call from the parent company's CEO, confirming the payment
- The victim said that he recognized "his boss' voice and slight German accent" as well as the "cadence and careful enunciation", so he made the payment to the scammers

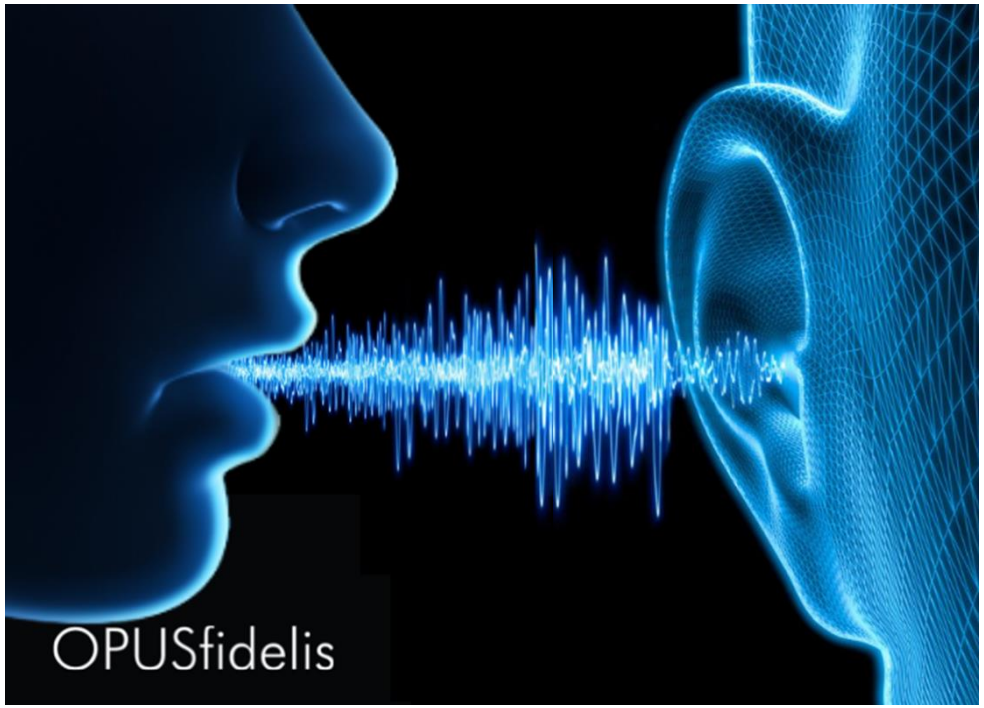

Malicious File Attachments

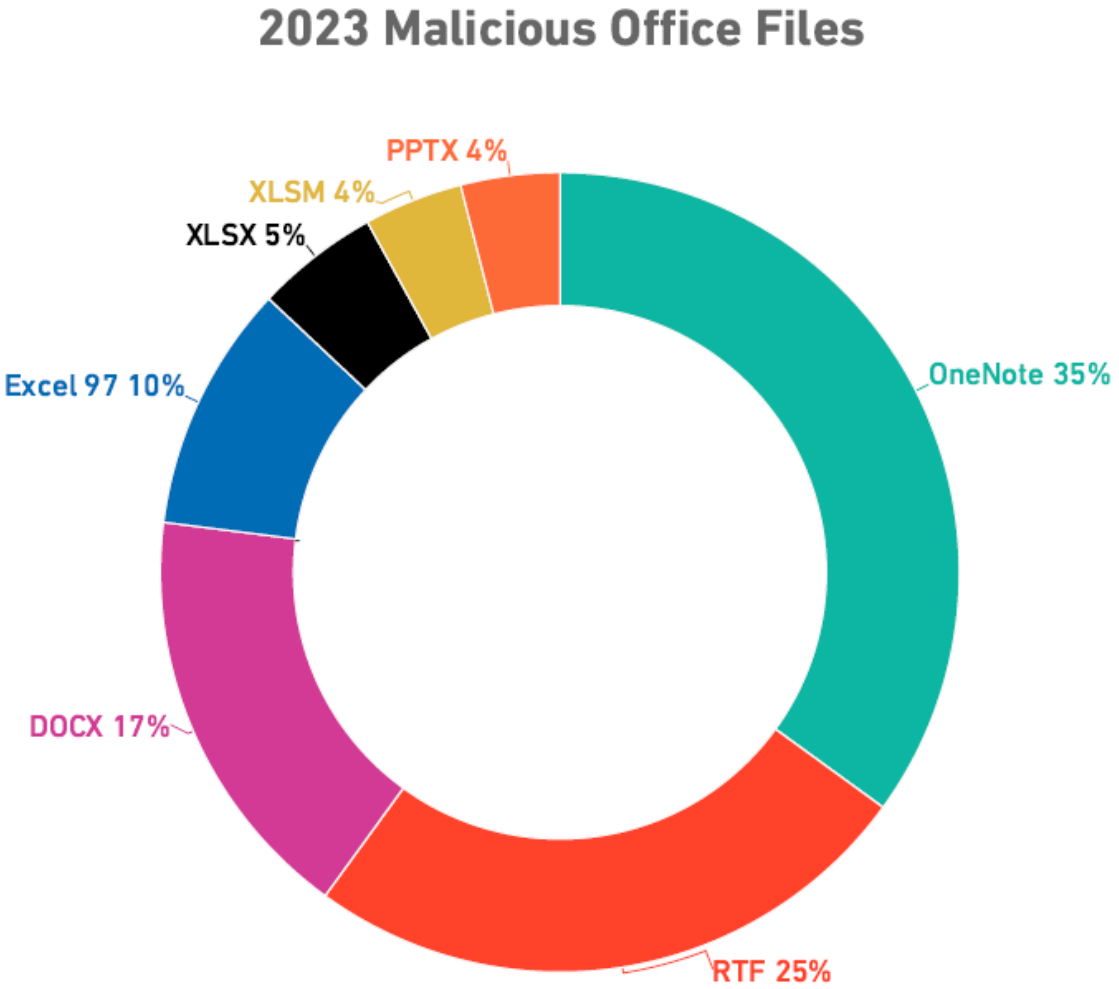

Remains the Top Attack Vector

Incident Response In A Nutshell

- **Contain** – Disconnect systems from network, disable switches in broad-base attack, disable outbound traffic
- **Eradicate** – Delete malware, disable affected user accounts, patch or mitigate vulnerabilities exploited, determine attack vector and implement controls against future attacks
- **Recover** – Restoration of files and data from secure and reliable backup sources
- **Post Incident** – identify root cause and create plans to prevent future incidents, identify communication and process flaws, evaluate response procedures and staff security training, and enforce security polices

Source: Microsoft Security, *Ransomware Incident Response Playbook Framework*

Ransomware Response

1. Determine which systems are impacted, and immediately isolate them
2. If you are unable to disconnect devices from the network, power them down to avoid further spread of the ransomware infection
3. Triage impacted systems for restoration and recovery, with mission critical systems actioned first
4. Consult with your incident response team to develop and document an initial understanding of what occurred based on their initial analysis
5. Engage your internal and external teams and stakeholders with an understanding of what they can provide to help you mitigate, respond to, and recover from the incident
6. Take a system image and memory capture of a sample of affected devices, such as workstations and servers
7. Contact local and federal law enforcement

Source: www.cisa.gov/shields-up

Summary

- If you thought technology had moved fast in the past, hang on!
- AI, Quantum and other emerging technologies will impact you
- Applications continue to proliferate, but less may be more
- Technology complexity is so great, few comprehend the risks