Highlights

WEKA 4 First Multicloud Data Platform for AI & Next-Gen Workloads

Harvard Graduate School of Design 2022 Wheelwright Prize & Data Storage

TikTok moves US Traffic to Oracle Servers- Buzzfeed contends that this step will not prevent Chinese access to US data

The Future of Western Digital and SanDisk

Google Employee Calculates Pi to 100 Trillion Digits & Sets New Record

Upcoming Conferences

WEKA 4 First Multicloud Data Platform for AI & Next-Gen Workloads

WEKA, the Data Platform for AI company, unveiled the fourth generation of its unified software-based platform that delivers a consistent, highly performant and scalable data management experience across on-premises, edge, hybrid – and now, multi-cloud – environments, for organizations looking to innovate and achieve first-to-market results with artificial intelligence (AI), machine learning (ML) and other next-generation workloads.

A seismic shift to the cloud by organizations looking to embrace digital-first strategies and harness transformational insights and discoveries using AI and ML is now accelerating, due in large part to the global disruptions wrought by the COVID-19 pandemic and subsequent supply chain constraints.

According to Flexera's 2022 State of the Cloud Report, eighty-nine percent of respondents reported having a multi-cloud strategy. There are several advantages to the multi-cloud architecture for innovative and/or large businesses. For one thing, it allows companies to easily move their data from one cloud to another - they aren't locked in to any one cloud provider. This flexibility makes companies more data resilient. Multi-cloud architectures also allow companies to specialize their workloads by cloud service.

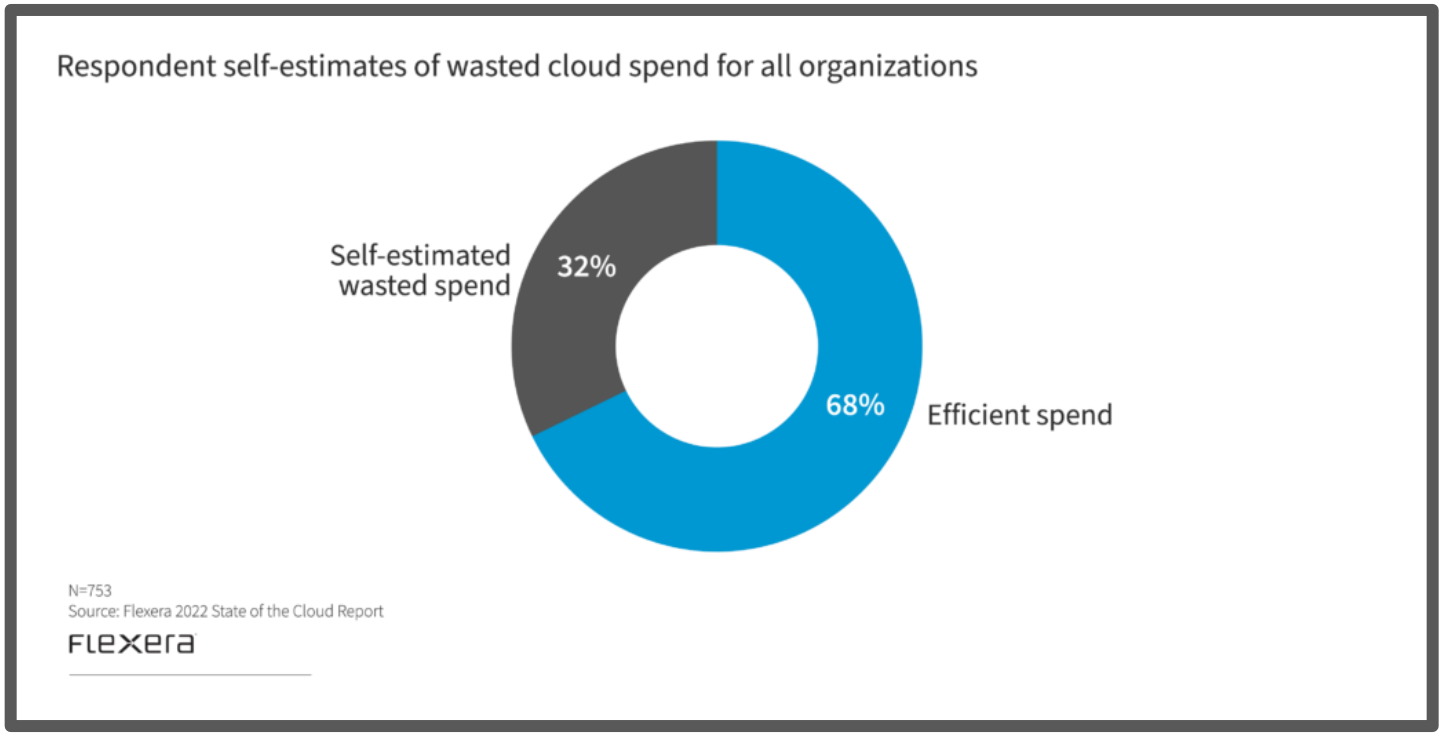

As AI and ML grow, the need for highly specialized software and hardware architectures to maximize AI and ML capabilities has become readily apparent. Advanced computing systems need the equivalent of high octane fuel to run optimally, whereas basic computing tasks like email and document creation only need to fill up on regular unleaded fuel. Segregating these systems allows companies to find scalable cost-savings in using different computing architectures. What is special about Weka 4 is that it allows companies to access all of these different systems through one single platform. Having everything accessible via one platform makes it much easier to find cost savings throughout the cloud.

For the sixth year in a row, optimizing use of cloud computing for cloud savings is the number one initiative for organizations. While firms are clearly very interested in optimizing their cloud usage, it appears to be quite the difficult task for organizations to complete! WEKA 4 provides three overlapping ways for customers to optimize their environments for cost, availability, scale – tiering to local and cloud object storage, new choice of capacity (QLC) or performance-optimized (TLC) NVMe drives, and new, filesystem-wide data reduction that can increase effective capacity.

"Hyperscale public clouds like AWS, Azure, GCP and OCI can provide the requisite agility and economies of scale needed to fuel these critical transformation and innovation engines, but the WEKA Data Platform is the key to unlocking that value for AI/ML workloads in hybrid and now multi-cloud environments. WEKA can uniquely help enterprises to avoid cloud lock-in and run their businesses with unparalleled economics."

Liran Zvibel, CEO at WEKA

Harvard Graduate School of Design 2022 Wheelwright Prize & Data Storage

Marina Otero's winning proposal, "Future Storage: Architectures to Host the Metaverse", explores a more socially and ecologically conscious architectural model for storing data. As data centers are growing in size and complexity, it's important to consider ways in which architectural design can reduce data centers' environmental footprint. The $100K Wheelwright Grant funds two years of research for the architect, who plans to travel to Iceland and Sweden for site visits as both countries are global leaders in renewable energy.

"While her title sounds futuristic, the issue is anything but: Marina recognizes the very urgent challenge of storing data in and for today's world. Her research will highlight innovations in data storage architectures and infrastructures, recognizing current inequities and scarcities, but also the potential for how data can transform entire communities worldwide through these new civic infrastructures and their reach. In her capacity as head of the social design masters at the Design Academy Eindhoven and also as director of research at the Het Nieuwe Instituut, she has already laid the groundwork for this important topic, ensuring the impact of her Wheelwright output, which will result in the first manual for global data center architecture design, as well as open-source course material and public programming," says Sarah M. Whiting, Harvard GSD's Dean and Josep Lluís Sert Professor of Architecture.

The (Wheelwright) prize reaffirms my confidence in the ability of this research to bring about new paradigms for consuming and storing data, expressly to make a difference. Data centers might not seem like an exciting place for an architectural project. However, the huge scale of the operations of the data industry and its pervasiveness and increasing importance in the contemporary world–coupled with its openness to innovation and concurrent pressures to find better socio-ecological models–creates a fertile environment for experimentation and action. - Marina Otero

TikTok moves US Traffic to Oracle Servers

Buzzfeed contends that this step will not prevent Chinese access to US data

TikTok recently announced that they shifted all of their US traffic to be routed through Oracle Cloud Infrastructure instead of their personal data centers in the United States and Singapore. While TikTok still uses their personal data centers as a backup, they expect to fully pivot to Oracle cloud servers over time. According to TikTok, they are working with Oracle to develop data management protocols that Oracle will audit to provide an extra layer of security. Furthermore, TikTok worked with US-based leadership to create a new department to solely manage US user data for TikTok.

TikTok claims that these changes will enforce additional employee protections, provide more safeguards, and further minimize data transfer outside of the US. **However, it remains unclear if these measures will restrict Chinese access to US data at all.** An investigative BuzzFeed report broke on the same day as TikTok's press release (June 17th), in which leaked audio from 80 internal TikTok meetings demonstrated the vast scope of Chinese access to US user data. Here are some reactions:

"Physical location does not matter if the data can still be accessed from China." - Adam Segal, Chair in Emerging Technologies and Director of the Digital and Cyberspace Policy Program at the Council on Foreign Relations

"It's almost incorrect to call it Oracle Cloud, because they're just giving us bare metal, and then we're building our VMs [virtual machines] on top of it." - TikTok's Head of Global Cyber and Data Defense, Will Farrell.

"I feel like with these tools, there's some backdoor to access user data in almost all of them." - Booz Allen Hamilton consultant, who was involved with TikTok's data migration to Oracle.

The Future of Western Digital (and SanDisk)

In 2016 Western Digital officially acquired SanDisk for approximately $16B. The original deal was announced at the end of 2015, and it came on the heels of the EMC Dell deal. Western Digital's acquisition of SanDisk was designed to integrate the two companies into the storage space, with SanDisk giving Western Digital market share in the growing SSD market. The transaction led to very little market consolidation as the two companies' product offerings were so different. However, it gave Western Digital a piece of the growing SSD space as its HDD business was slowly shrinking, and this diversification ensured that the company would stay relevant in the storage space even if HDD became obsolete.

With Dell's purchase of EMC, EMC's products benefitted from Dell's position as a leading hardware vendor. Dell could resell EMC's products (and EMC could resell Dell's products), and they were able to combine software and hardware in ways that made their combined offerings complimentary. In contrast, the Western Digital and SanDisk deal combined two hardware manufacturers building different components with different long term goals. Remember, the hard disk drive business had/has been shrinking as the SSD business has been growing as a percentage of total storage sales. While SSDs were/are taking over the consumer storage market, HHDs were/are fighting to hold onto that market share. So conflict was inevitable between the two groups once SanDisk became a part of Western Digital, especially when it comes to long-term R&D investments and capital outlays.

Now Western Digital is the target of activist shareholder group Elliott Investment Management, who owns a 6% share of the company. Elliot wants Western Digital to spin off the SSD portion of the business from HDDs and inject $1B of capital into the SSD business. They say that such a move could significantly raise the value of the company, which has underperformed the NASDAQ over the last six years since the merger occurred. However, general investor sentiment is skittish at this point, and the stock had dropped from $60.08 to $46.35 since Elliott wrote the letter to the board at the beginning of the month (of course, the market has taken a dive in general, so how much to attribute to this situation versus the doom and gloom of the day, is not clear).

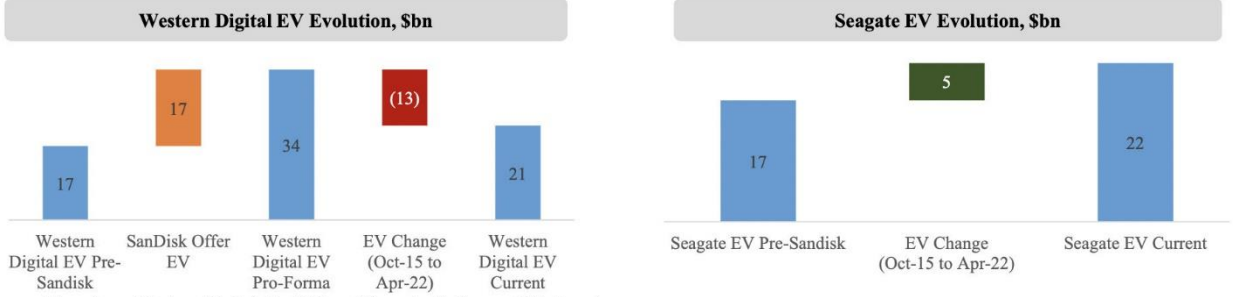

Source: Bloomberg, Western Digital, SanDisk and Seagate disclosure, Elliott analyses.

"We believe a full separation of the Flash business can allow both HDD and Flash to be more successful and unlock significant value. By executing on a separation, we believe Western Digital's stock price could reach $100+ per share by the end of 2023, representing uniquely attractive upside of approximately 100%". - Elliott Investment Management

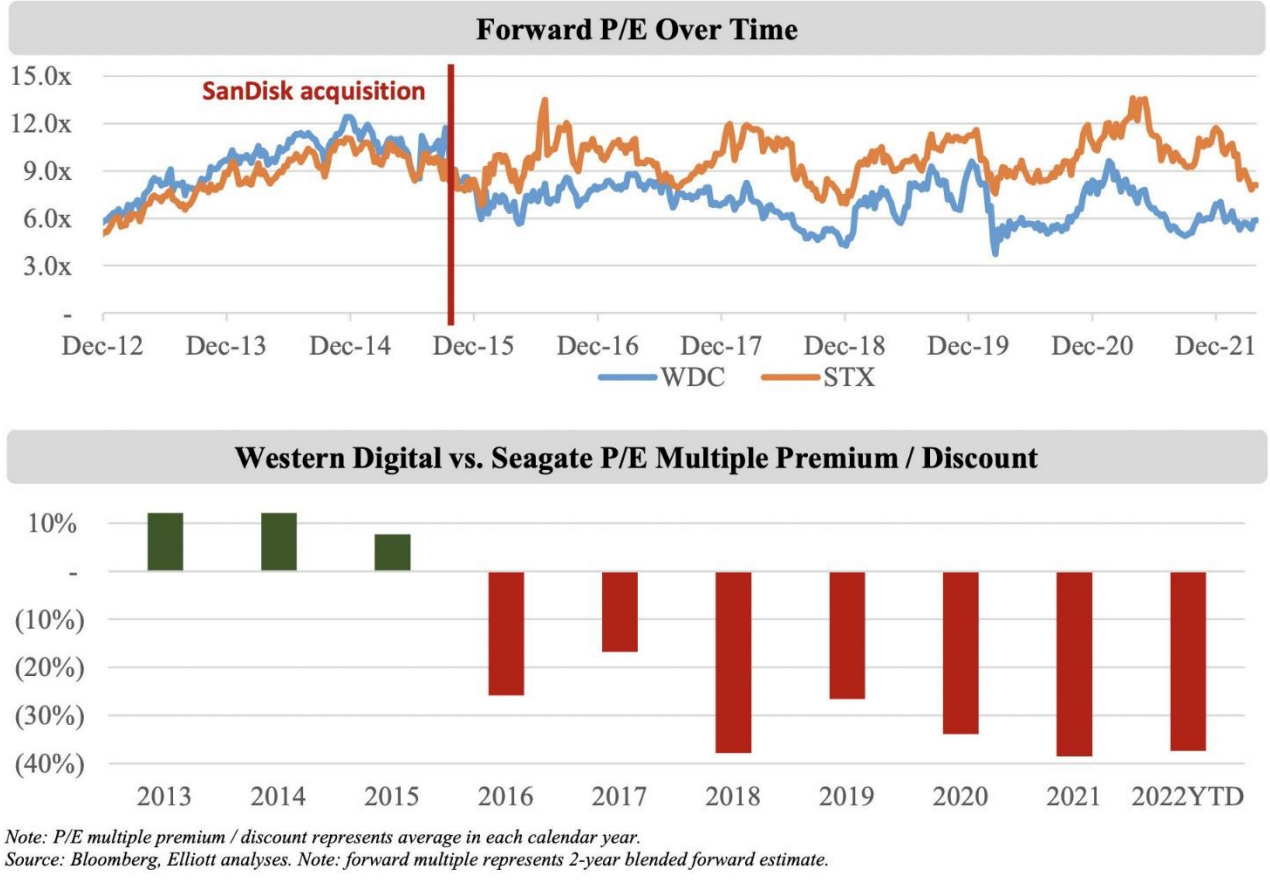

Note: P/E multiple premium / discount represents average in each calendar year.
Source: Bloomberg, Elliott analyses. Note: forward multiple represents 2-year blended forward estimate.

Interestingly enough, Elliott also targeted EMC in 2014. At the time, Elliot was a 2% shareholder in EMC, and suggested that the EMC "federation" with VMware was obscuring VMware's value, and that EMC should either spin off VMware or make it part of a M&A with another company to maximize shareholder value. EMC was unhappy with this, as well as all of the public scrutiny that comes with being a publicly listed company. Thus EMC was very happy to be taken private by Dell in what was the largest pure tech deal in history.

Ironically Dell went public again a few years later (2018), ostensibly to help it pay back the large sums of debt that it had accumulated from the EMC deal. Furthermore, Dell ended up spinning off VMware in 2021 as Elliott suggested EMC should do in 2014. While tech companies tend to hate the public oversight that activist investors bring to their company, it's also clear that M&As in the tech space are fraught with challenges to overcome. While these larger companies in theory can exert more market

pressure on competitors, and may be able to cross sell their products to each other's client base, there are significant organizational challenges for acquiring companies to overcome. This includes workplace culture, long term goals, management structure, and organizational competence in successfully managing acquired companies. Certainly Western Digital and SanDisk have performed worse than they did pre-merger, as they have lost market share in both HDDs and SSDs.

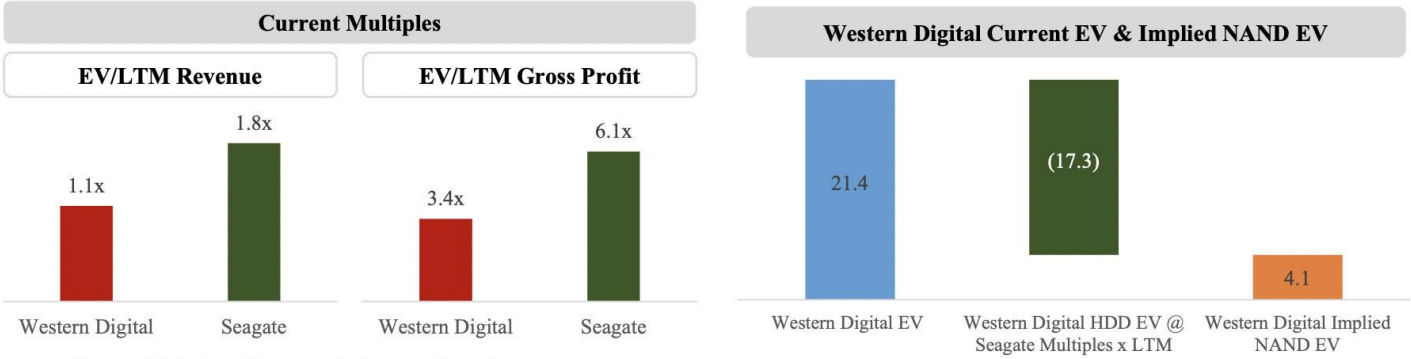

Source: Western Digital and Seagate disclosure, Bloomberg.

Given all of this, what does the future hold for Western Digital? Will the division of the SSD's from the HDD's lead to better management of both products? Elliott certainly believes in the SSD portion of Western Digital, as evidenced by their willingness to inject $1B of capital into the SSD spinoff. Certainly there is a lot of growth in the SSD portion of the storage market, with a compounded annual growth rate between 10.2% and 17.2%. Western Digital is significantly undervalued on an earning basis, trading at an 8.02 P/E ratio (as of 6/16/22).

Part of the challenge will be rebuilding the culture and leadership that SanDisk had. As an activist investor group, Elliott does not have the same power that a private equity firm has in installing its own management and/or board of directors. Still, plenty of opportunity exists in a spinoff. It will be interesting to see what happens to Western Digital and SanDisk, and the possibility of spinning off SanDisk to KIOXIA is also very interesting, as both groups have a joint partnership to design and build 3D NAND flash memory. Western Digital states that it will consider splitting the HDD and SSD portions of its business.

The Board is aligned in the belief that maximizing value creation warrants a comprehensive assessment of strategic alternatives focused on structural options for the company's Flash and HDD businesses. Through this process, we are actively engaging in a broad range of strategic and financial alternatives that will help further optimize the value of Western Digital, including Elliott's offer to invest incremental equity capital in our Flash Business. We look forward to continuing our constructive dialogue with Elliott as this process unfolds.

Western Digital CEO David Goeckeler

Google Employee Calculates Pi to 100 Trillion Digits & Sets New Record

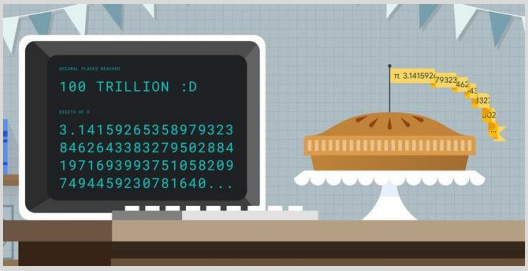

Emma Haruka Iwao set a new world record for her calculation of pi to 100 trillion digits with Google. This calculation surpassed her previous record of 31.4 trillion digits that was set in 2019. Her calculation took 157 days to calculate, from October 2021 to March of 2022. Considering that it took her team 121 days to do a 31.4 trillion digit calculation in 2019, this is a more than double increase in calculation speed on a per digit basis.

In many ways the feat shows off the computing power of Google Cloud. The team used balanced Persistent Disks, a cost-effective SSD solution on a per-GB basis that offers 1,200 MB/s read and write throughput and 15-80k IOPS.

Here are more of the technical details of the setup used to run the calculation:

Program: y-cruncher v0.7.8, by Alexander J. Yee
Algorithm: Chudnovsky algorithm
Compute node: n2-highmem-128 with 128 vCPUs and 864 GB RAM
Start time: Thu Oct 14 04:45:44 2021 UTC
End time: Mon Mar 21 04:16:52 2022 UTC
Total elapsed time: 157 days, 23 hours, 31 minutes and 7.651 seconds
Total storage size: 663 TB available, 515 TB used
Total I/O: 43.5 PB read, 38.5 PB written, 82 PB total

Meaning that 5.15 bytes of storage were used per digit calculated, 385 bytes of data were written per digit calculated, and 435 bytes of data were read per digit calculated!

Emma on Breaking the PI Record Twice!

Breaking the record of π was my childhood dream, so a few years ago I decided to try using Google Cloud to take on this project. I also wanted to see how much data processing these computers could handle. In 2019, I became the third woman to break this world record, with a π calculation of 31.4 trillion digits.

But I couldn't stop there, and I decided to try again. And now we have a new record of 100 trillion decimal places. This shows us, again, just how far computers have come: In three years, the computers have calculated three times as many numbers. - Emma Haruka Iwao

If you read all 100 trillion digits out loud, one second at a time, it would take you 3,170,979 years to read the whole thing.

Emma and LGBTQ+ Community Pride

As an openly queer woman and software engineer, visibility is important for me and my community.

Being recognised in this way by Guinness is such an honour, and I can only hope that it will serve as a source of inspiration for the LGBTQ+ community and beyond. - Emma Haruka Iwao

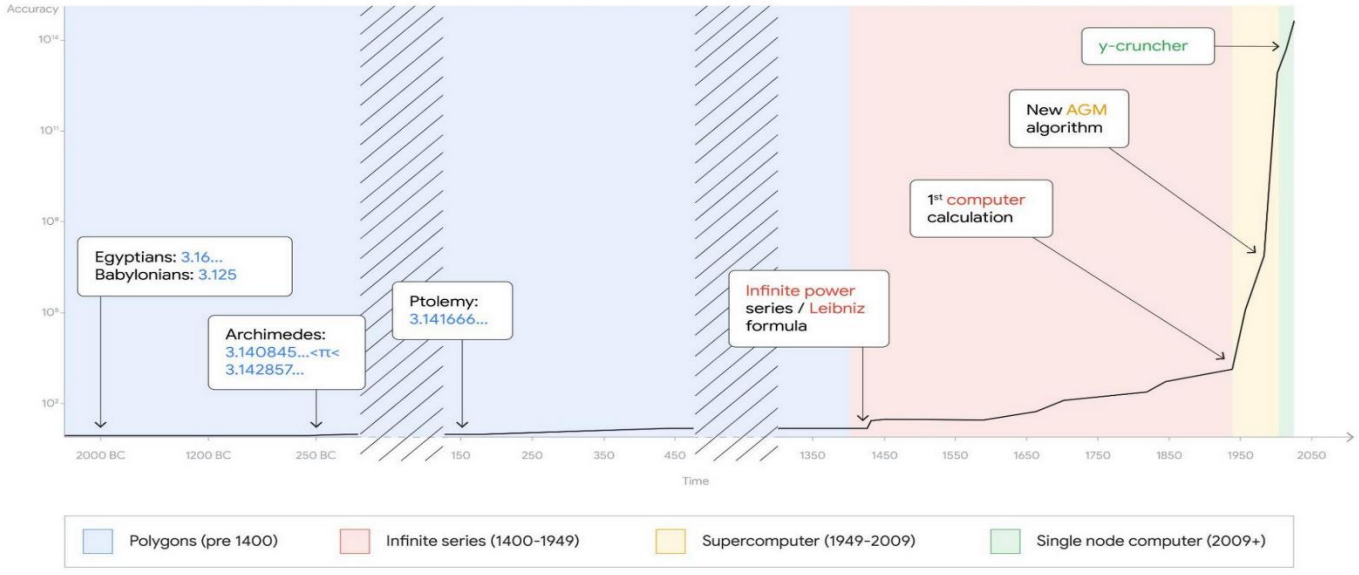

How Big is 100 Trillion?

The height of a stack of 100,000,000,000,000 (one hundred trillion) one dollar bills measures 6,786,616 miles. This would reach from the Earth to the Moon and back 14 times!

If you were to count that stack of one dollar bills at a rate of one bill per second, it would take you 3.2 million years to count the entire stack!

100 trillion inches of pie crust would stretch from Earth to the Moon and back 3,304 times!

Upcoming Conferences

June 21-22	Gartner Security & Risk Management Summit, Sydney
June 21-22	Gartner Digital Workplace Summit, San Diego
June 29- July1	Mobile World Congress, Shanghai
July 13	FutureCon St. Louis, Hybrid
July 19-20	Cyber Solutions Summit & Expo, Virtual
July 25-27	Gartner Security & Risk Management Summit, Tokyo
August 2-4	Flash Memory Summit, Santa Clara
August 6-11	Black Hat USA, Vegas
August 11-14	DEF CON 30, Vegas
August 27-28	Blue Team Con, Chicago
August 29-Sept 1	VMwear Explore, San Francisco
September 8	FutureCon, Des Moines
September 12-14	Gartner Security & Risk Management Summit, London
September 13-14	CISO Forum, Virtual
September 14	Cybersecurity Expo, Phoenix
September 19-20	Industry of Things World, Berlin
September 20-22	Dreamforce, San Francisco
September 22-23	Global Cyber Conference, Zurich
September 26-28	InfoSec World, Colorado Springs
September 27-28	International Cyber Expo, London

September 28-29	IoT World, Santa Clara
September 28-30	Spiceworld, Austin, Hybrid
October 3-4	451Nexus, Las Vegas
October 5-6	Evolve, Vegas
October 6-7	Big Data & AI Toronto
October 10-12	ISC Security Congress, Las Vegas
October 11-13	Google Cloud Next, Virtual
October 17-19	Authenticate 2022, Seattle
October 17-20	Gartner IT Symposium/Xpo, Orlando
October 24-27	ICS Cybersecurity Conference, Hybrid/Virtual
November 16	San Diego Cybersecurity Conference, Hybrid
November 16	Threat Hunting Summit, Virtual
November 18-19	Data Strategy & Insights (Forrester Research), Virtual
November 21-22	Gartner IT Infrastructure, Operations, & Cloud, London
November 28-Dec 2	AWS re:Invent, Las Vegas
December 1-2	AI & Big Data Expo Global, London
December 6	Security Operations Summit, Virtual