Artificial Intelligence Machine Learning and Deep Learning for the Network Engineers

https://t.me/learningnets

1. Introduction to AI in Networking

Definition of AI, Machine Learning (ML), and Deep Learning (DL)
Why AI is relevant to Network Engineers
Key AI drivers in networking: automation, optimization, and analytics
Overview of AI applications in networking:

- Network Design
- Network Operations
- Network Security

2. Foundations of AI, ML, and Deep Learning

Fundamental AI Concepts, Terminologies, and Applications
Types of Machine Learning:

- Supervised, Unsupervised, and Reinforcement Learning

Deep Learning basics and its relevance to networking

agenda

3. AI/Machine Learning in Network Design

3.1 Business Needs

Aligning AI/ML implementations with organizational goals

Identifying networking challenges that AI can address

3.2 Data Sovereignty

Managing data across public, private, and hybrid cloud environments

Compliance with regional data regulations

3.3 Security

Ensuring secure implementation of AI/ML in networking

Addressing vulnerabilities in AI-driven systems

3.4 Assurance

Guaranteeing reliability and availability of AI-integrated networks

Managing AI failures or anomalies

3.5 Integrity

Ensuring data accuracy and consistency in AI processes

Validation methods for AI-driven decisions

3.6 Impacts

- Storage requirements and traffic pattern changes due to AI adoption
- Infrastructure and resource optimization strategies

3.7 Auto Scalability

- Leveraging AI for dynamic scalability in networking

3.8 Cost and ROI

- Evaluating the financial impact of AI integration
- Measuring ROI for AI-driven network solutions

3.9 Governance

- Establishing policies for AI use in networking
- Ethical considerations and operational frameworks

3.10 Sustainability

Green AI
Affordability(CPU/GPU tweaking)
Effective use of AI accelerators
Power and cooling requirements

3.11 AI Network Design Use Cases

Machine learning for predictive network modeling
Large Language Models for network design insights and automation
Pattern recognition for identifying network bottlenecks and optimization opportunities
Impacts on infrastructure resources and requirements for different AI/ML use-cases
a Machine learning
b Deep learning
c LLM
d GenAI

4. Networking for AI

High-Performance Networking Technologies

This section focuses on technologies designed to enable low-latency, high-bandwidth communication for AI workloads.

RDMA
InfiniBand
RoCE
RoCEv2
iWARP
Ultra Ethernet
UEC

Compute Accelerators and Interconnects

This section delves into the hardware and interconnects that power AI computation and facilitate seamless communication between compute nodes.

DPU, TPU, GPU
NVLink
NVMe-over-Fabrics (NVMe-oF)

4. Networking for AI

Traffic Management and Congestion Control

This section focuses on techniques and technologies used to manage and optimize network traffic for AI workloads, ensuring reliable and efficient communication.

Layer 2

DCB – Datacenter Bridging

- **PFC (Priority Flow Control)**
- **ETS (Enhanced Transmission Selection)**
- **QCN (Quantized Congestion Notification)**
- **DCBX (Datacenter Bridging Exchange)**

Layer 3

ECN (Explicit Congestion Notification)

DCQCN (Data Center Quantized Congestion Notification)

TIMELY

HPCC (High Precision Congestion Control)

Introduction to AI in Networking

- Definition of AI, Machine Learning (ML), and Deep Learning (DL)
- Why AI is relevant to Network Engineers
- Key AI drivers in networking: automation, optimization, and analytics
- Overview of AI applications in networking:
    - Network Design
    - Network Operations
    - Network Security

agenda

What is AI , ML and DL

- Artificial intelligence (AI) makes it possible for machines to learn from experience, adjust to new inputs and perform human-like tasks

- Machine Learning and Deep Learning are subset of Artificial Intelligence

What is AI , ML and DL

Artificial Intelligence

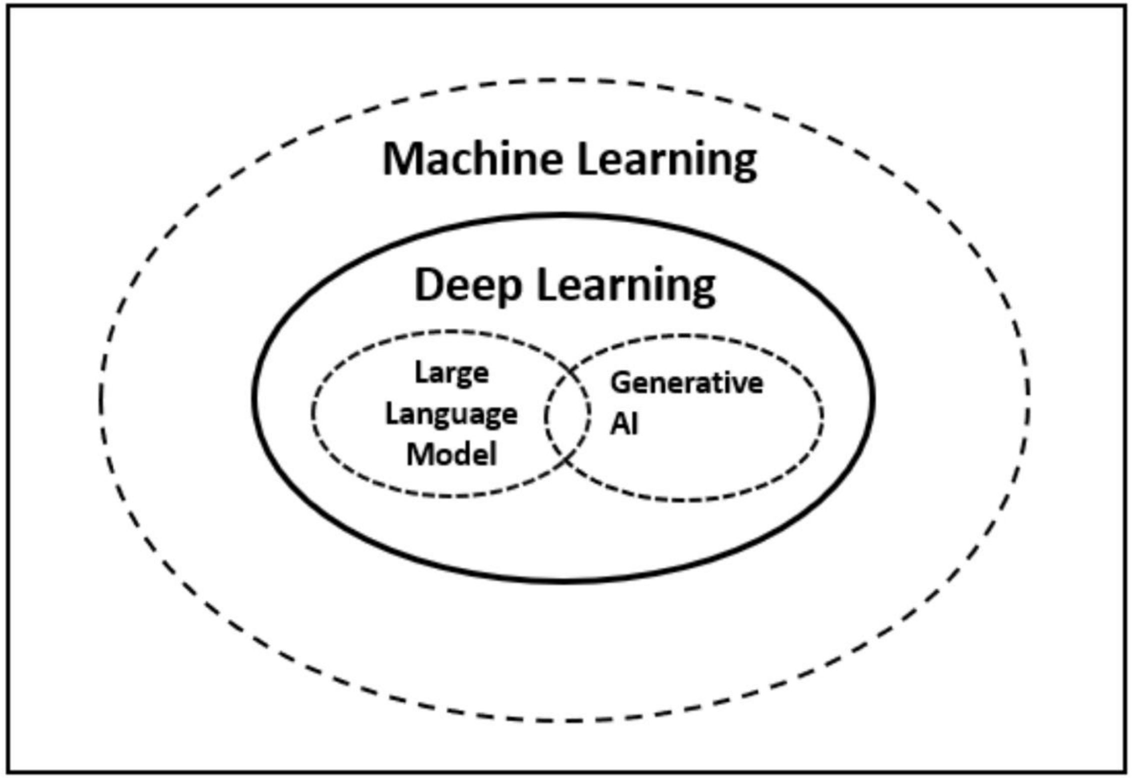

https://t.me/learningnets

What is AI , ML and DL

Artificial Intelligence (AI)

- **Simple Definition:** AI is like teaching computers to think and act like humans

- **Example Use:** Virtual assistants like **Alexa** or **Siri** understand your voice and give answers

What is AI , ML and DL

Machine Learning (ML)

ML is when computers learn from data to make decisions without being directly programmed

Netflix suggests movies you might like based on what you've already watched.

What is AI , ML and DL

Deep Learning (DL)

- **Simple Definition:** DL is a type of ML that uses brain-like systems called neural networks to solve very complex problems

- **Example Use: Self-driving cars** recognize stop signs and other cars using deep learning

Why AI is relevant to Network Engineers?

Automation of Repetitive Tasks:

AI can automate routine tasks like configuration management, fault detection, and performance monitoring, allowing network engineers to focus on strategic activities

Enhanced Network Security:

AI-powered anomaly detection and threat intelligence systems proactively identify and mitigate security threats in real-time

Performance Optimization:

AI helps optimize routing, bandwidth allocation, and QoS policies by analyzing traffic patterns and predicting future network demands

Why AI is relevant to Network Engineers?

Proactive Issue Resolution:

Predictive maintenance and failure analysis reduce downtime by identifying potential issues before they impact the network

Improved Decision-Making:

AI-powered analytics provide actionable insights for network design, capacity planning, and SLA compliance

Handling Big Data:

Networks generate vast amounts of telemetry data. AI efficiently processes and extracts insights from this data for better network visibility and control

Key AI drivers in networking: automation, optimization, and analytics

Automation:

•AI automates repetitive and time-consuming tasks, such as configuration management, fault detection, and patch deployment.

•Enables **zero-touch provisioning** and **intent-based networking**, reducing human errors and improving operational efficiency.

•Example: Automatically configuring network policies in SD-WAN environments based on application needs.

Key AI drivers in networking: automation, optimization, and analytics

Optimization:

•AI dynamically optimizes network resources, such as bandwidth allocation, traffic routing, and Quality of Service (QoS)

•Helps reduce latency, improve throughput, and maximize resource utilization, even in complex multi-cloud and hybrid environments

•Example: AI-based traffic engineering to reroute data during congestion in real-time

Key AI drivers in networking: automation, optimization, and analytics

Analytics:

•AI processes and analyzes large volumes of telemetry and log data to provide actionable insights

•Supports predictive maintenance, anomaly detection, and capacity planning by identifying trends and potential issues

•Example: Using AI to detect unusual traffic patterns indicative of a potential DDoS attack.

Overview of AI applications in networking:

- Network Design
- Network Operations
- Network Security

Overview of AI applications in networking:

Network Design

•**Role of AI:**

- Assists in designing efficient and scalable networks by automating complex tasks like topology optimization and capacity planning.
- Predicts future network requirements based on historical data, ensuring proactive design decisions.

•**Applications:**

- AI-powered tools generate intent-based designs aligned with business objectives.
- Traffic simulation and predictive modeling to validate network resilience under varying conditions.

Overview of AI applications in networking:

Network Design Use Cases:

- **Cisco:**
  - **Cisco DNA Center:**
    - Uses AI to provide **intent-based network design**, helping organizations align their network topology and policies with business objectives.
    - Example: AI-driven insights recommend optimal device placement and configurations for SD-Access fabrics.

- **Other Vendors:**
  - **Juniper Networks' Apstra:**
    - Automates the design and lifecycle management of data center networks using AI-based intent validation and topology optimization.
    - Example: Validates underlay and overlay configurations to ensure consistent network design.

Overview of AI applications in networking:

Network Operations

•**Role of AI:**

- Streamlines daily operations with intelligent automation and real-time monitoring.
- Reduces downtime through predictive maintenance and proactive fault detection.

•**Applications:**

- Self-healing networks using AI to identify and fix issues automatically.
- Dynamic traffic engineering to optimize resource utilization and minimize congestion.

Overview of AI applications in networking:

Network Operations

- **Cisco:**
  - **Cisco AI Network Analytics (part of DNA Center):**
    - Uses AI to monitor network performance in real time, detect anomalies, and provide actionable insights.
    - Example: Predicts potential bandwidth congestion and automatically adjusts traffic paths to maintain optimal performance.

- **Other Vendors:**
  - **Arista Networks' CloudVision:**
    - AI-powered network telemetry and automation platform for monitoring and troubleshooting.
    - Example: Identifies misconfigurations in EVPN-VXLAN networks and suggests corrective actions.

Overview of AI applications in networking:

AI Use Cases In Network Monitoring

By using AI capabilities, businesses can enhance their network monitoring practices in various ways. Here are some use cases of AI in network monitoring:

- **Anomaly detection**: AI network monitoring tools can quickly identify unusual patterns or deviations from normal network behavior, which might indicate a security breach or system failure.
- **Predictive analytics**: By analyzing historical data, AI can predict potential network failures or performance degradations before they occur.
- **Automated configuration and optimization**: AI can automate routine network configuration tasks and optimize network settings based on current traffic patterns and demands.
- **Security enhancement**: AI enhances network security by detecting and responding to threats in real time. It can identify malware, ransomware, and other malicious activities quickly, minimizing potential damage.
- **Root cause analysis**: When problems occur, AI can help diagnose the root cause more quickly than traditional methods. By correlating various data points and identifying patterns, AI reduces the time needed to troubleshoot and resolve issues.
- **Capacity planning**: AI can forecast future network needs based on trend analysis, helping organizations plan upgrades and expansions more effectively.

Overview of AI applications in networking:

Network Security

•**Role of AI:**

- Enhances security by detecting and mitigating threats faster than traditional methods.
- Leverages advanced anomaly detection to identify suspicious patterns in network traffic.

•**Applications:**

- Real-time intrusion detection and prevention.
- AI-driven threat intelligence to predict and counteract emerging attack vectors.

Overview of AI applications in networking:

Network Security

- **Cisco:**
  - **Cisco Secure Network Analytics (formerly Stealthwatch):**
    - Uses AI and machine learning to detect threats such as lateral movement and data exfiltration within the network.
    - Example: AI identifies unusual east-west traffic in a data center, indicating potential malware activity.

- **Other Vendors:**
  - **Palo Alto Networks Cortex XDR:**
    - AI-based platform for extended detection and response (XDR) across network, endpoint, and cloud.
    - Example: Uses AI to correlate alerts from multiple data sources, identifying advanced persistent threats (APTs) before they escalate.

Overview of AI applications in networking:

Networking use cases often fall into specific categories:

Use Case	AI/ML Category	Example Algorithms or Techniques
Anomaly Detection	Unsupervised ML	K-Means, DBSCAN, Autoencoders, Isolation Forest
Traffic Prediction	Supervised ML	Linear Regression, LSTMs, Transformers
Routing Optimization	Reinforcement Learning	Q-Learning, Deep Q-Networks, Actor-Critic
Traffic Classification	Supervised ML	Decision Trees, SVMs, Random Forests
Network Automation	Rule-Based AI	Expert Systems, Graph Neural Networks
Chatbots for Support	NLP	Transformers, BERT, GPT

Overview of AI applications in networking:

We talked about AI for Networking Use Cases

Network Security
Network Operations
Network Design
Predictive Maintenance
Network Monitoring

We will discuss Networking for AI in detail later in the course!

Foundations of AI, ML, and Deep Learning

Foundations of AI, ML, and Deep Learning

Fundamental AI Concepts, Terminologies, and Applications
Types of Machine Learning:
- Supervised Learning
- Unsupervised Learning
- Reinforcement Learning

Deep Learning basics and its relevance to networking

What is AI , ML and DL

ML and DL are subsets of AI

You perform AI using ML and DL technology

Differences between AI and ML

All ML solutions are AI solutions, but not all AI solution is ML

Key Differences

Aspect	AI	ML
Scope	Broad, encompassing ML and other methods like rule-based systems.	Focused on learning from data.
Learning Required?	May or may not involve learning from data.	Always involves learning from data.
Applications	General intelligence tasks like planning and reasoning.	Specific learning tasks like prediction and classification.
Examples	Virtual assistants, robotics, rule-based chatbots.	Recommendation engines, image recognition, anomaly detection.

Differences between AI and ML

AI is the big picture, encompassing any technology that mimics human intelligence, whether or not it uses data to learn

ML is a powerful tool within AI, designed to learn from data and improve decision-making or predictions automatically

Artificial Intelligence Applications

Voice Recognition

NETFLIX

Content Recommendation

Google Translate

Autonomous Driving

Artificial Intelligence Applications

Voice Recognition: AI powers assistants like Alexa or Siri to understand and respond to spoken commands

Content Recommendation: AI helps Netflix suggest shows and movies based on your viewing habits

Language Translation: AI enables real-time translation of text or speech across different languages (e.g., Google Translate)

Autonomous Driving: AI processes sensor data to help self-driving cars navigate safely and make decisions on the road

Drivers of Advances in AI

Processing

Algorithms

A Lot of Data

https://t.me/learningnets

Why AI Now? Key Drivers of Advances in AI

Abundance of Data:

•Massive amounts of data are generated daily (e.g., IoT, social media, and cloud computing), providing rich inputs for AI models to learn from

Processing Power:

•Advances in hardware, like GPUs, TPUs, and specialized AI accelerators, enable faster training and inference for complex AI models

Improved Algorithms:

•Innovations like deep learning, transformers, and reinforcement learning have drastically improved AI's ability to handle complex tasks with higher accuracy

Machine Learning

- Machine Learning is an algorithm that learn from data without explicitly programmed

- There are many types of Machine Learning but mainly all Machine Learning and Deep Learning algorithms are classified as either Supervised or Unsupervised Learning or Reinforcement Learning

Machine Learning vs. Traditional Programming

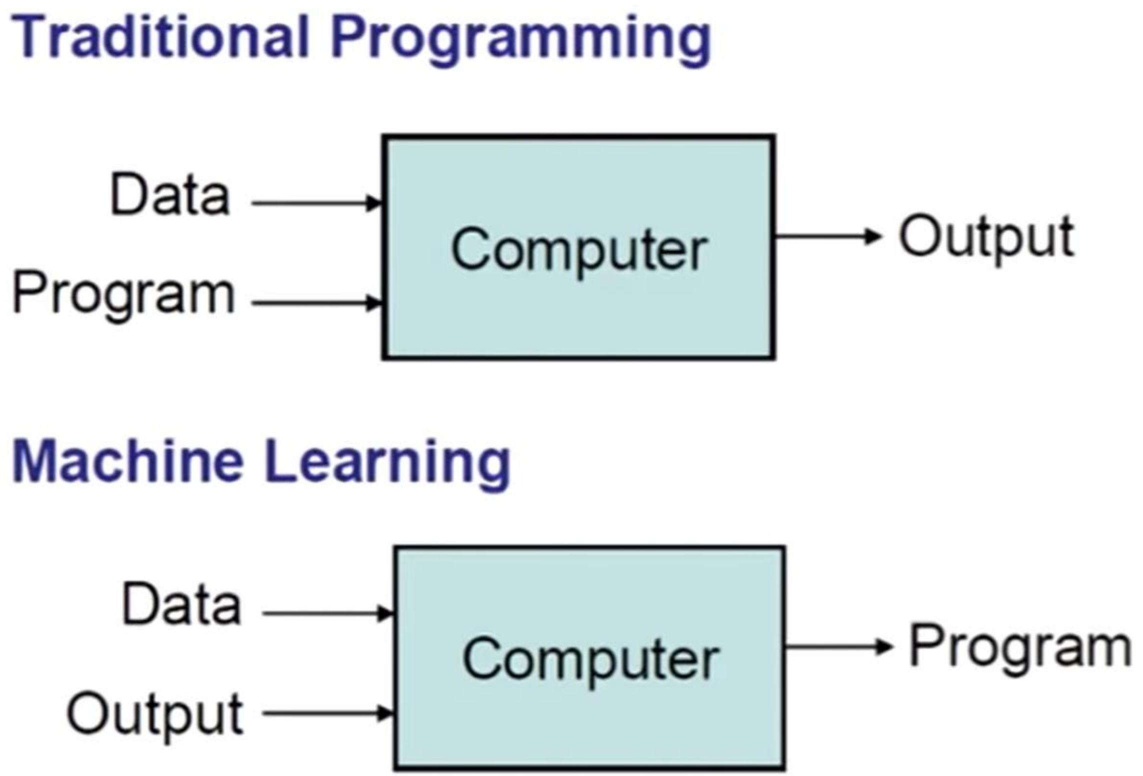

Machine Learning vs. Traditional Programming

Traditional Programming:

You write a program (rules and logic) and combine it with input data to get the output

Machine Learning:

You provide the **data** (inputs) and the **desired outputs**, and the system learns to create the program (the model) on its own

- **Example:**
  - For email spam detection:
    - Input: Email content.
    - Output: Spam or not spam.
    - The model learns the rules by analyzing patterns in the data

Machine Learning vs. Traditional Programming

In **traditional programming**, humans define the logic

In **machine learning**, the system discovers the logic by learning from the data and examples

Machine Learning Types

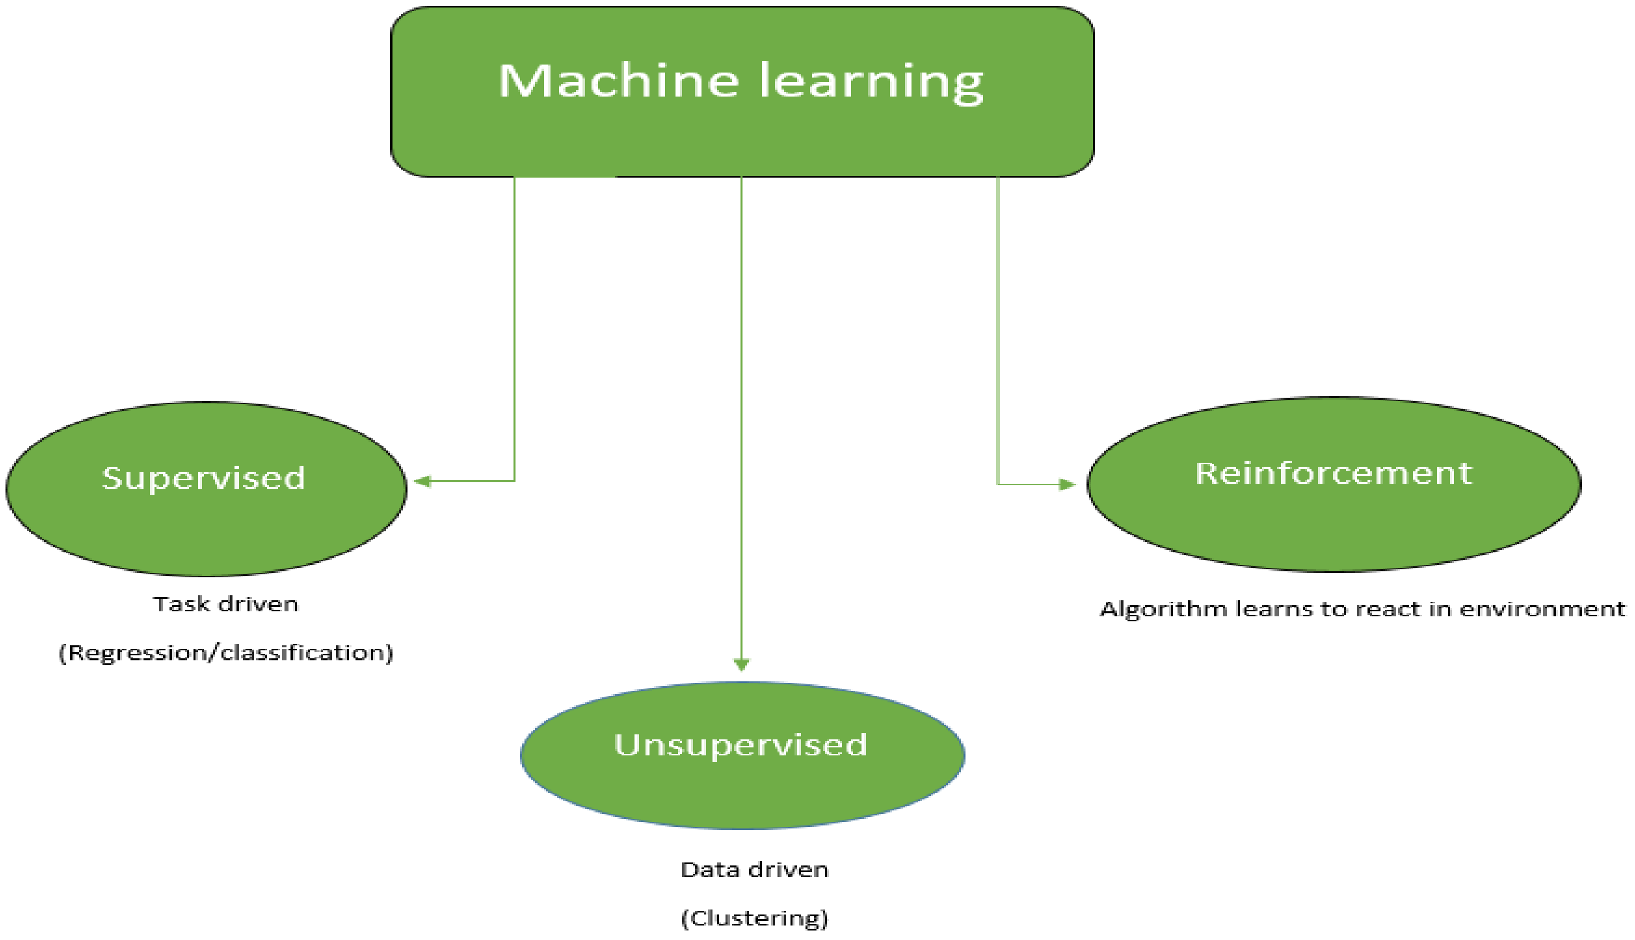

Machine Learning – Supervised Learning

Supervised learning uses labeled training data to train machines to learn relationships between given inputs to a given output

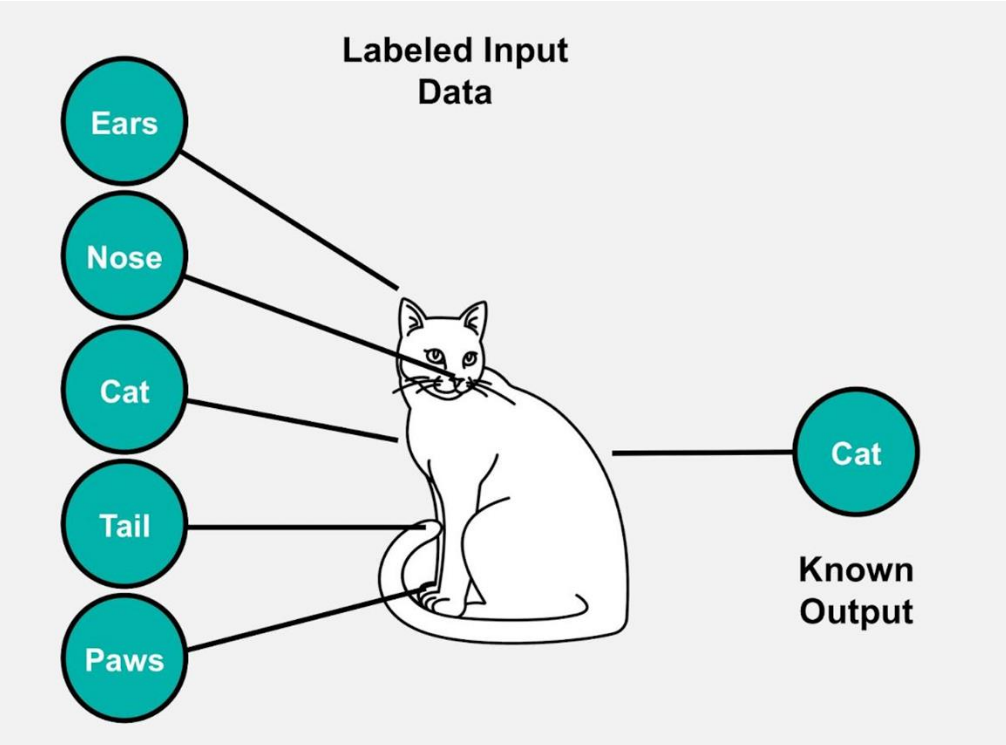

Machine Learning – Supervised Learning

Supervised Learning provides predictability for each labeled information

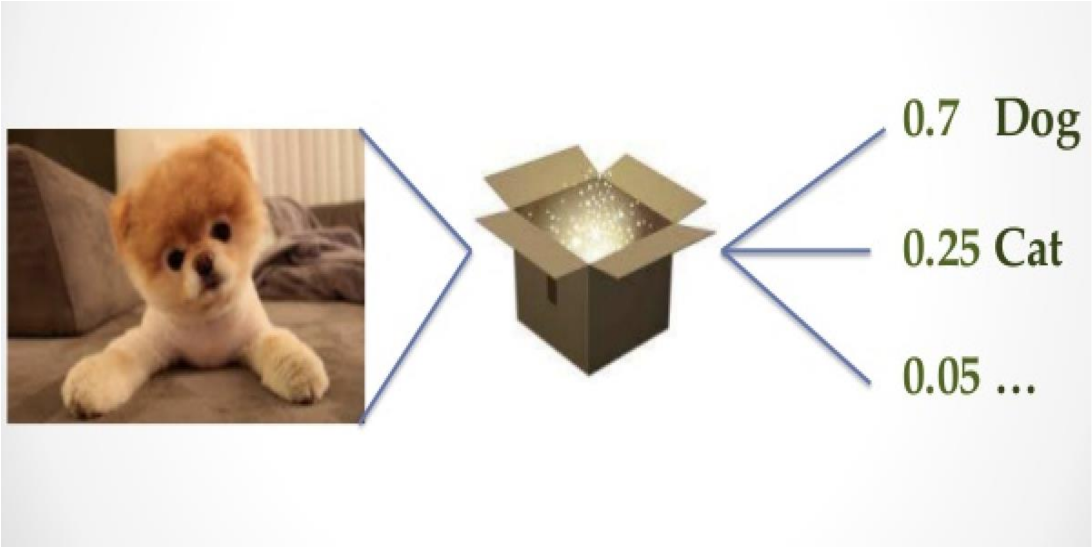

Supervised vs. Unsupervised Learning

- Supervised Learning separate each labeled data
- Unsupervised Learning cluster them

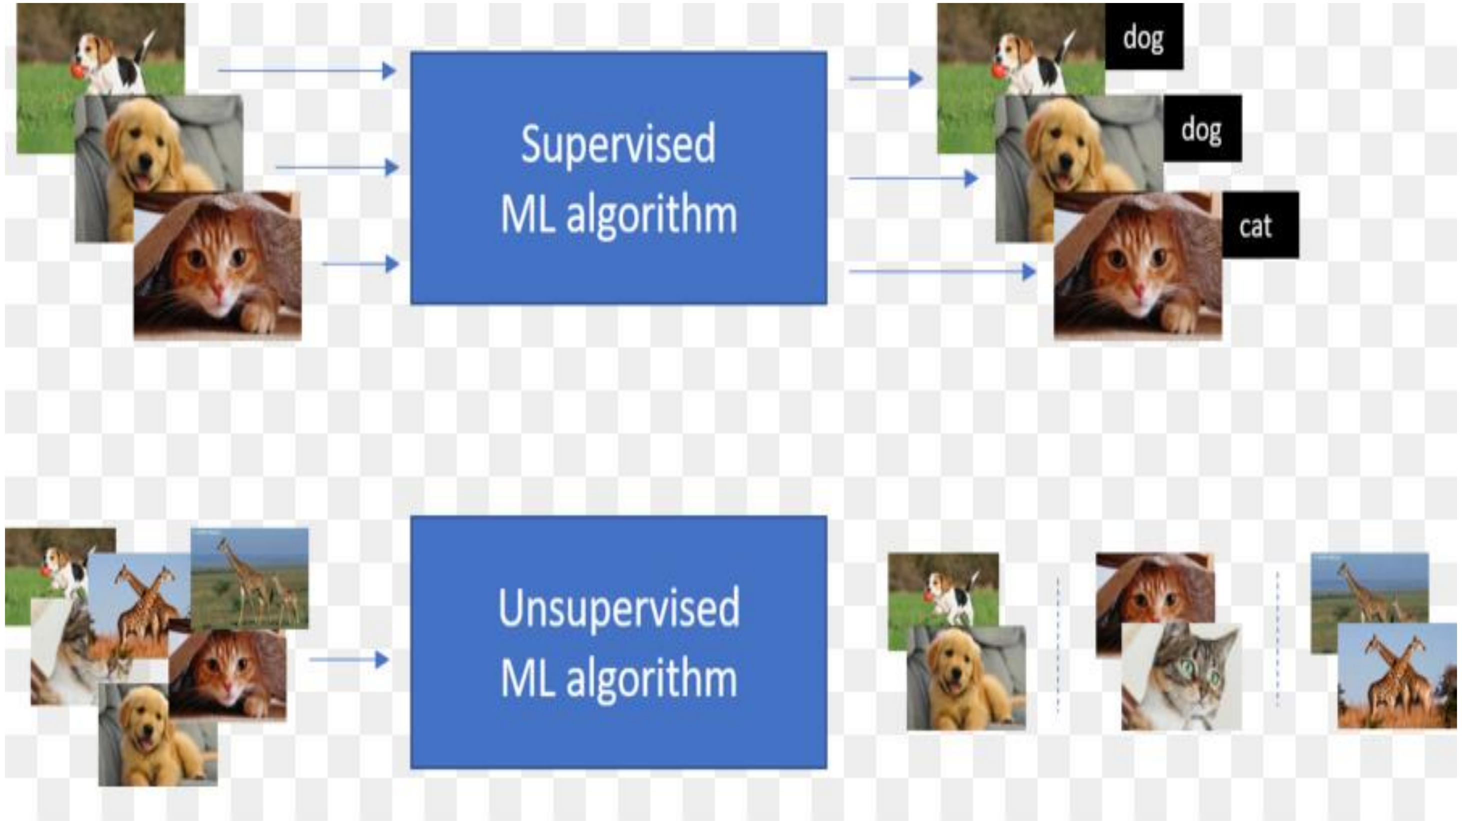

Machine Learning – Unsupervised Learning

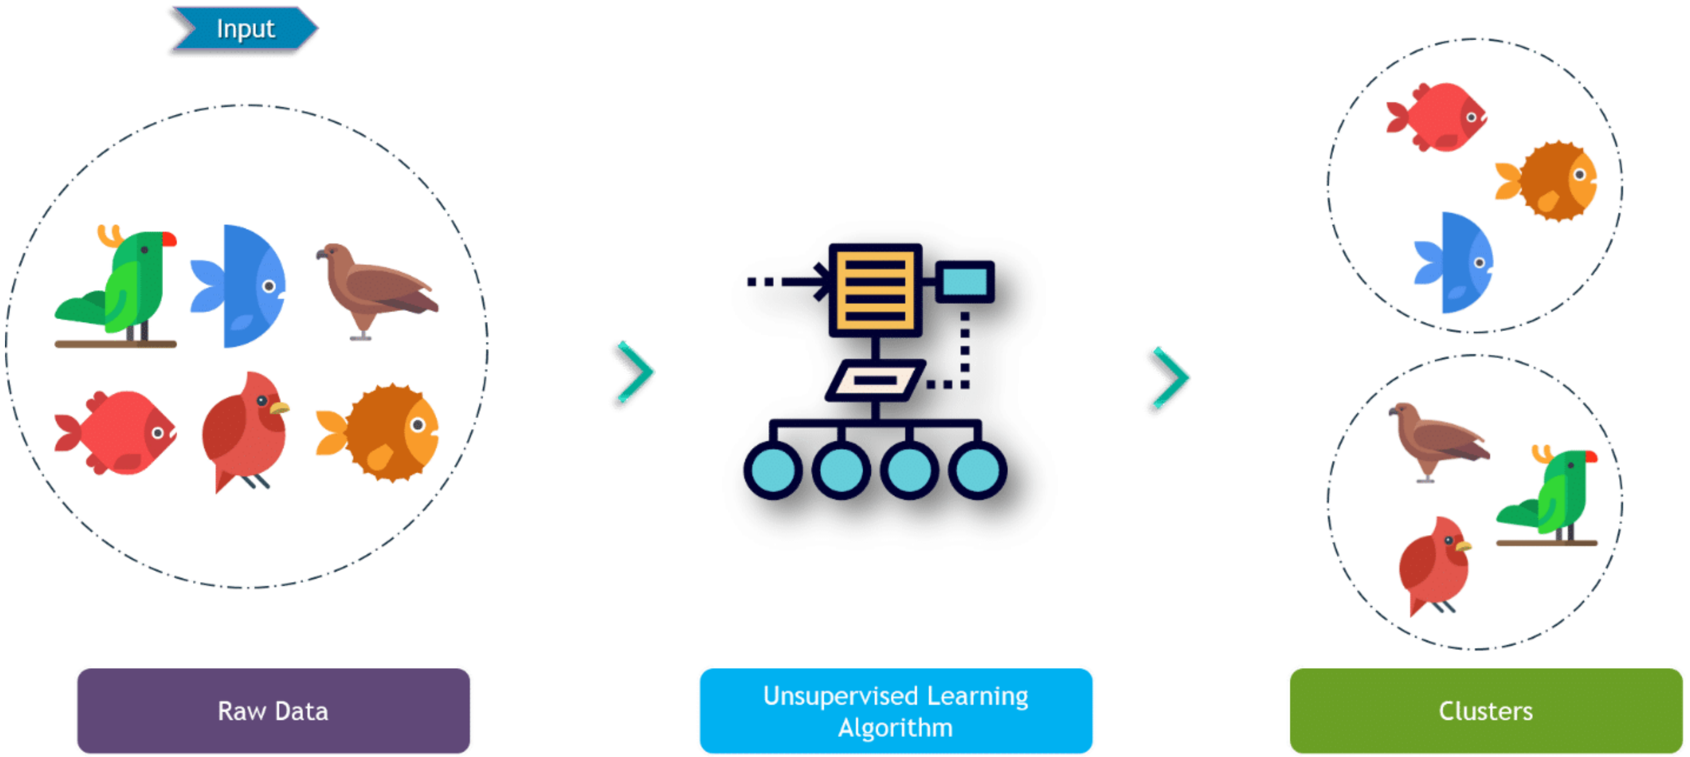

https://t.me/learningnets

Machine Learning – Reinforcement Learning

- Reinforcement learning (RL) is learning by interacting with an environment
- An RL agent learns from the consequences of its actions and it selects its actions on basis of its past experiences and also by new choices (exploration)
- Reinforcement Learning is essentially trial and error learning

Machine Learning – Reinforcement Learning

- The reinforcement signal that the RL-agent receives is a numerical reward, which encodes the success of an action's outcome, and the agent seeks to learn to select actions that maximize the accumulated reward over time

Machine Learning – Reinforcement Learning

- Dog training can be given as an example to Reinforcement Learning

- Whole meaning of reinforcement learning training is to “tune” the dog’s action so that it learns the desired behaviors by getting a reward when each time it performs correct action

Machine Learning – Reinforcement Learning

- After training is complete, the dog should be able to observe the owner and take the appropriate action, for example, sitting when command is to "sit"

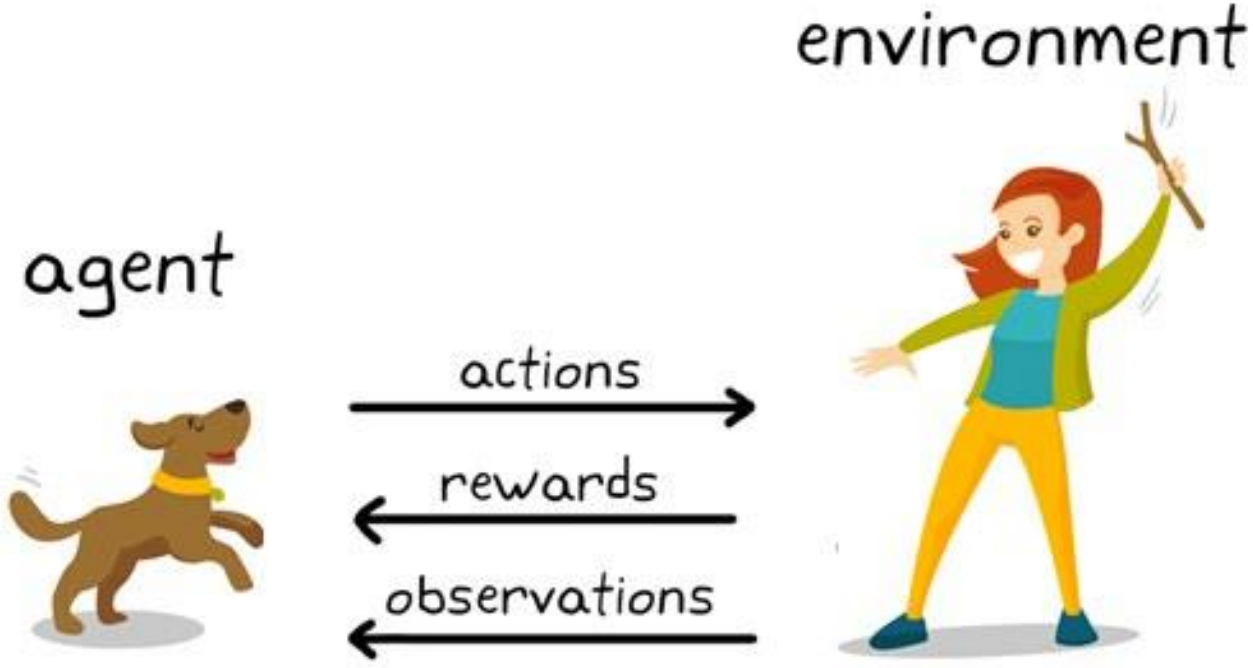

Real Life example of Reinforcement Learning

Autonomous Parking

Real Life example of Reinforcement Learning

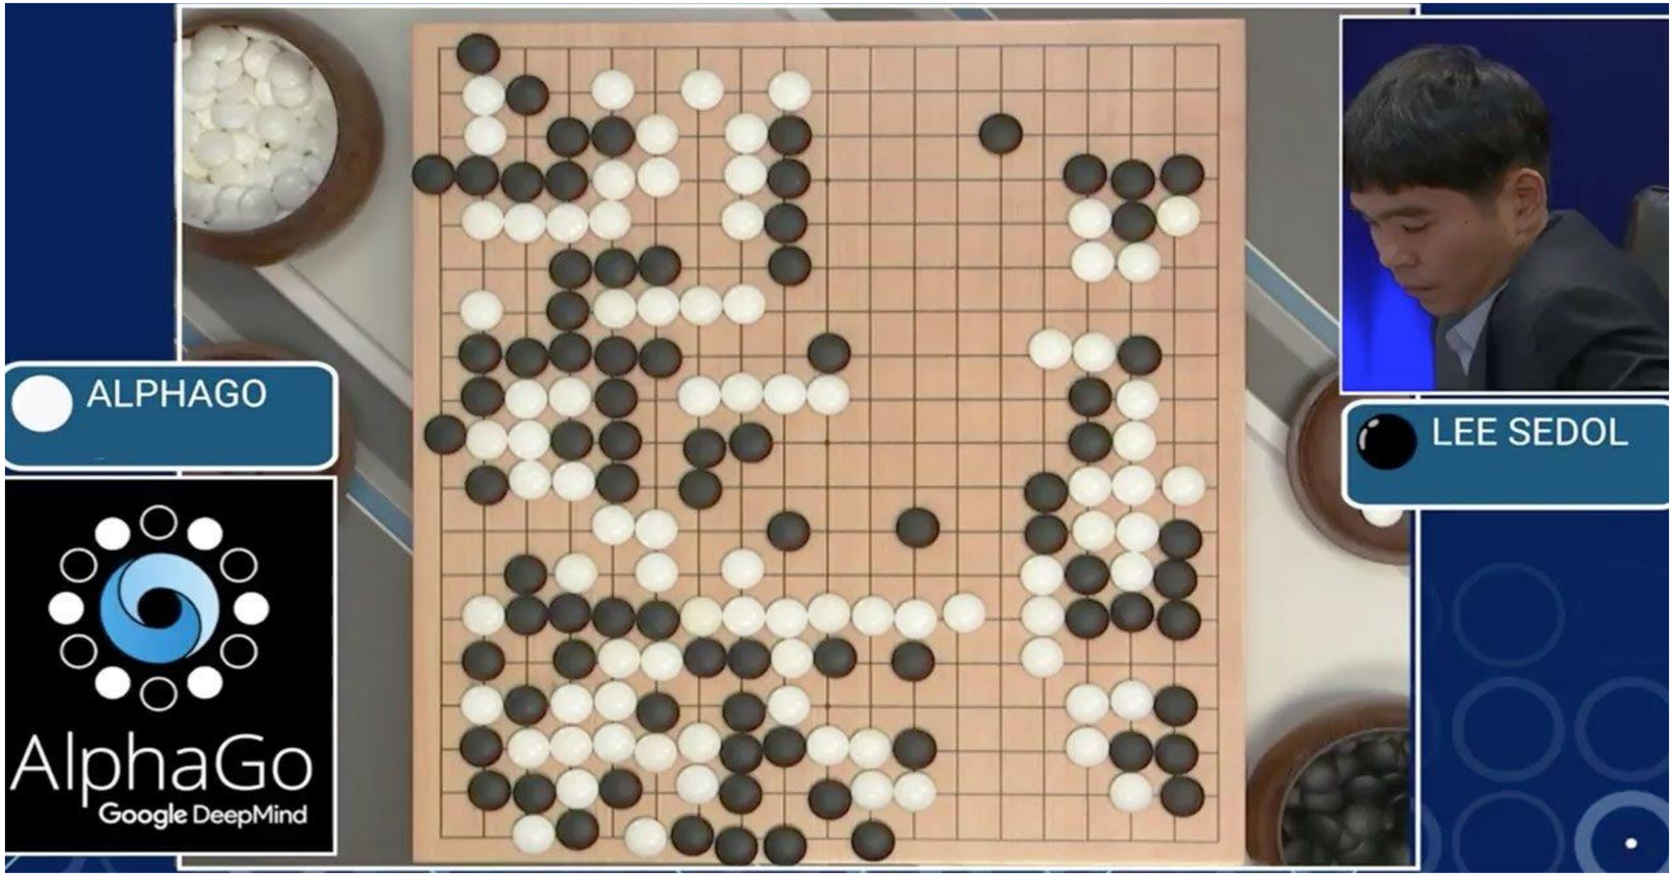

Real Life Examples of Supervised and Unsupervised Learning

Supervised and Unsupervised Learning are not just used to understand whether picture is cat or dog!

There are many real life applications which people can get benefit!

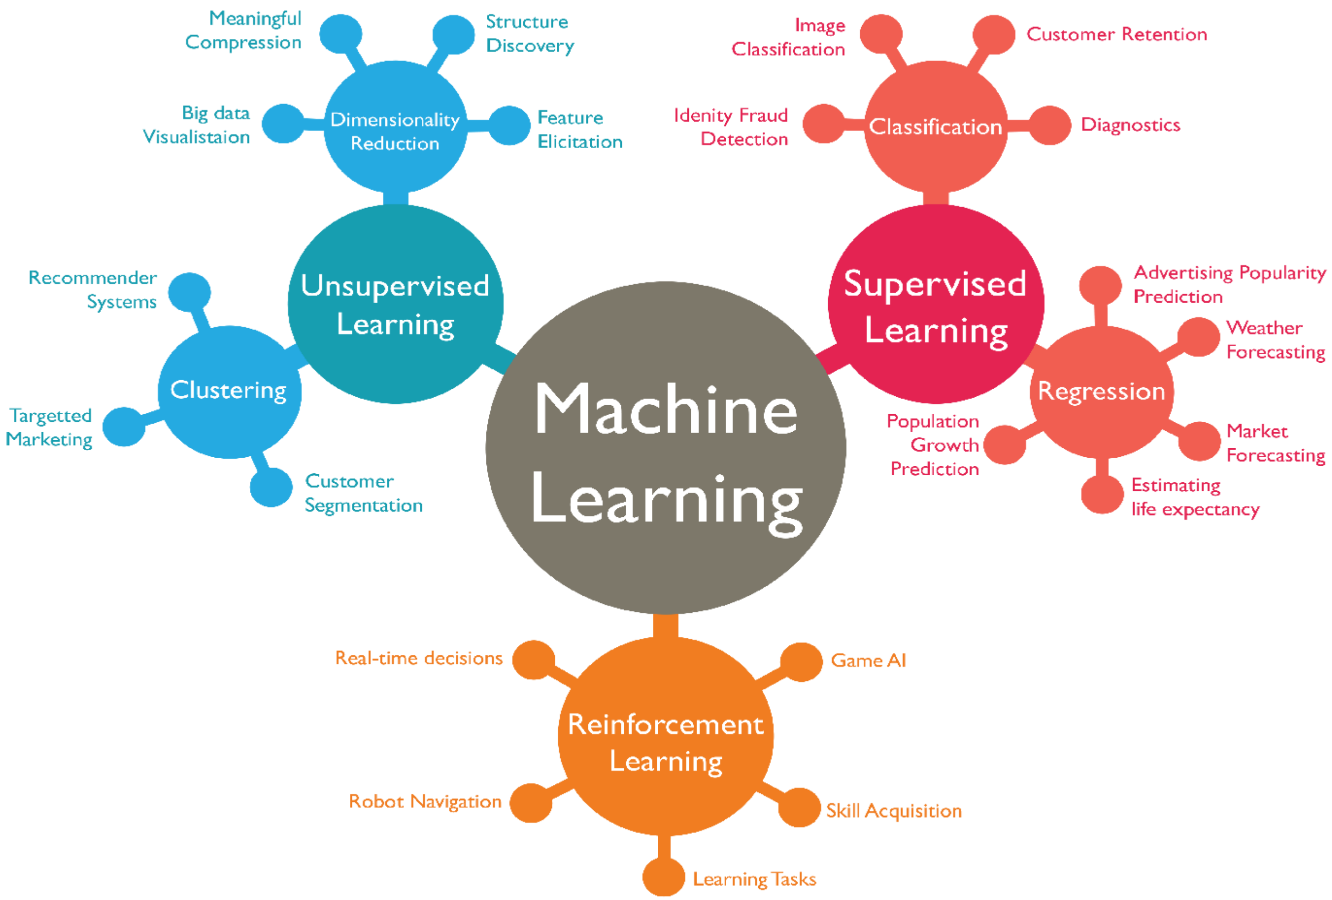

https://t.me/learningnets

Deep Learning (Neural Networks)

- **Deep learning is a machine learning technique**

- **Works Like the Human Brain**
  Deep learning uses neural networks, which mimic the way our brains process information, to solve complex problems.

- **Handles Complex Data**
  It's great for unstructured data like images, videos, and text, making it essential for tasks like facial recognition, language translation, and speech-to-text.

Deep Learning (Neural Networks)

- **Needs Big Data**
  Deep learning thrives when trained on massive datasets. The more data it has, the better it performs.

- **Resource Intensive**
  It requires significant computational power (think GPUs and TPUs) and time for training large models, which can be a limitation for some use cases.

Deep Learning (Neural Networks)

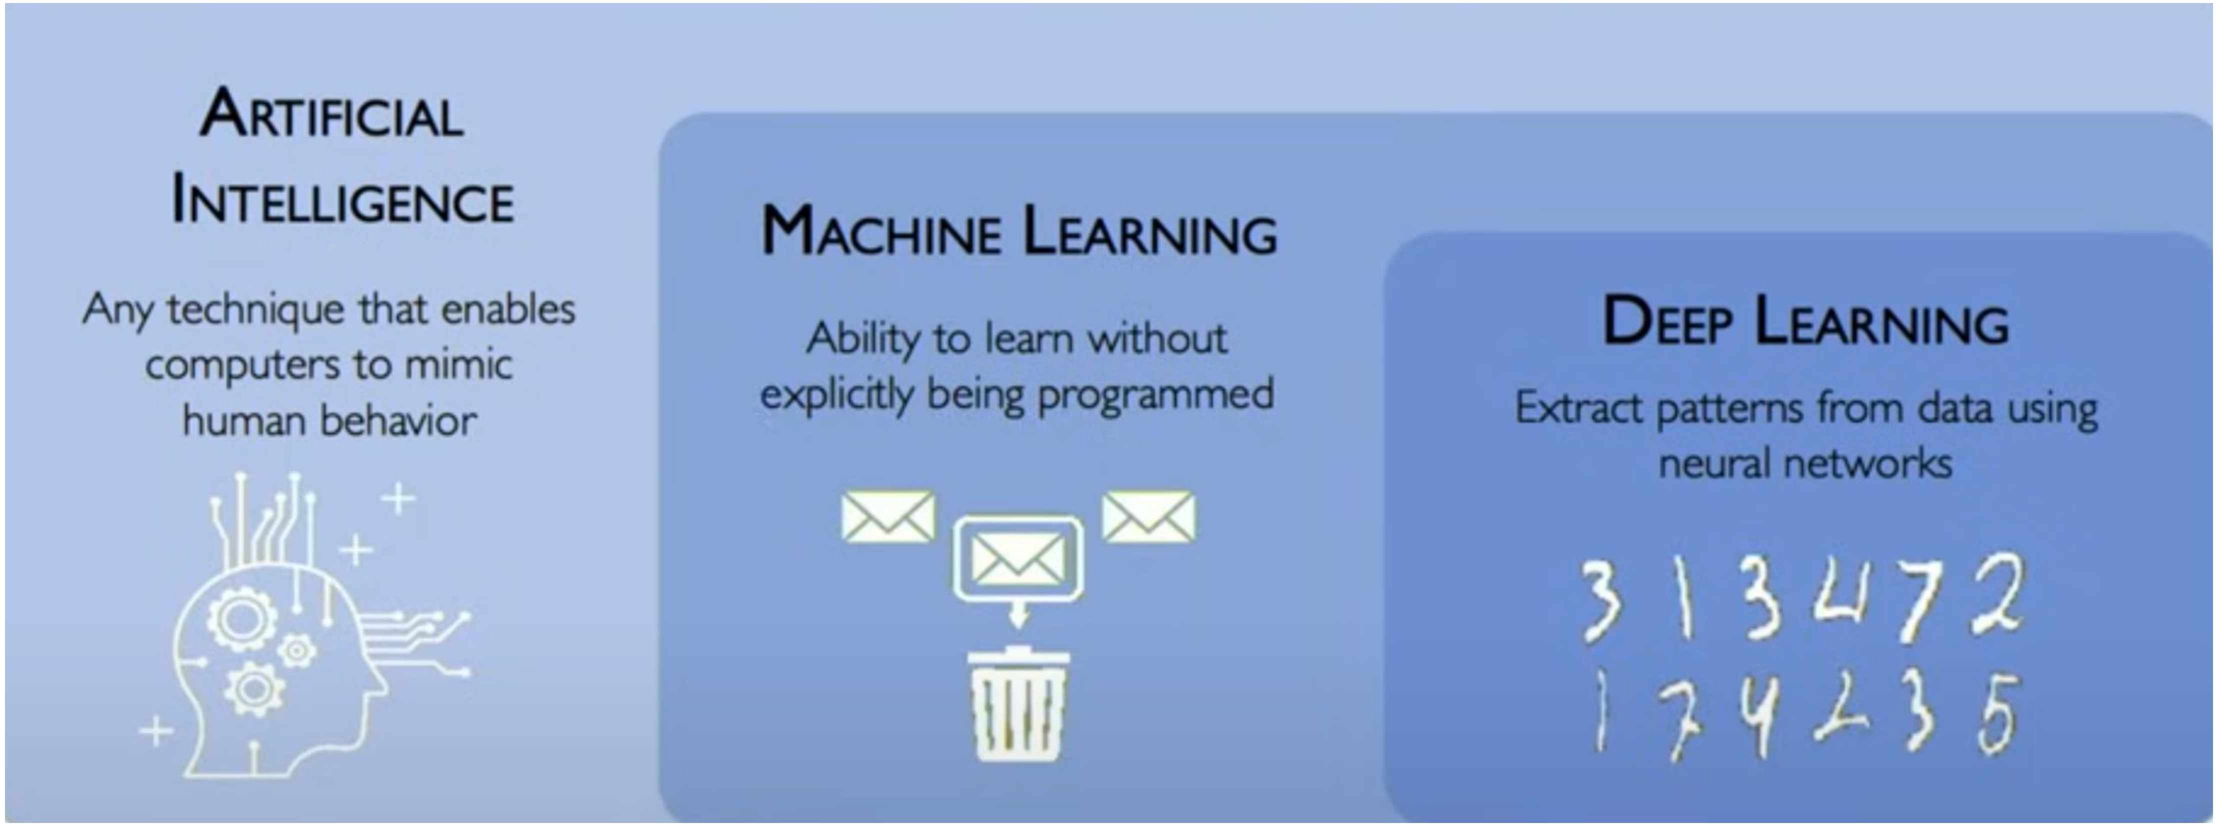

Deep Learning (Neural Networks) Applications

Facial recognition

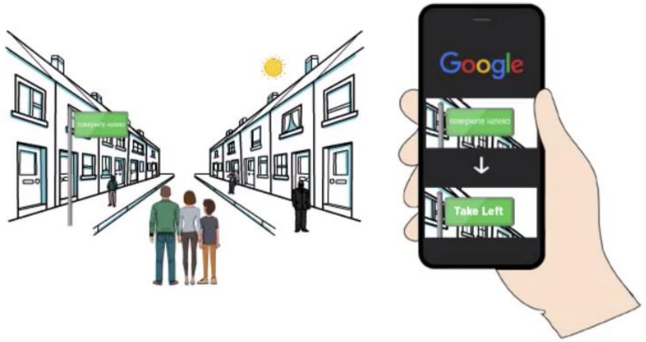

Real-time translation

Music composition

Deep Learning (Neural Networks)

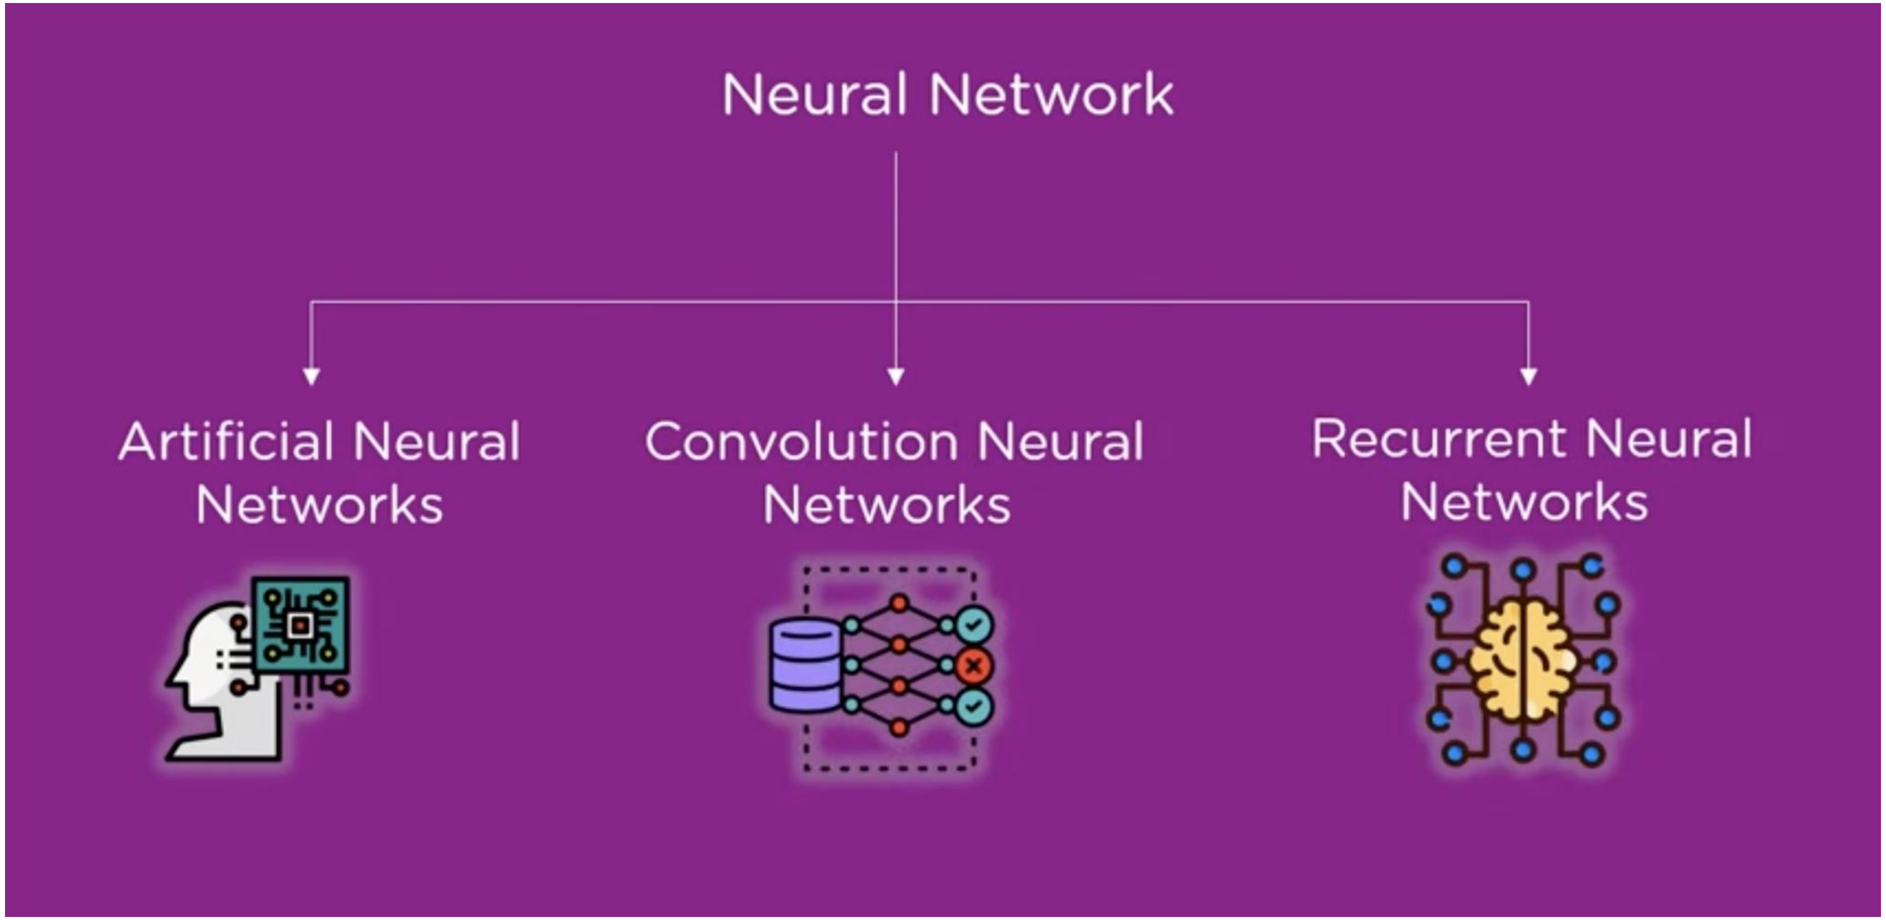

Deep Learning (Neural Networks)

- In the human brain, there are about 100 billion neurons. Each neuron connects to about 100.000 of its neighbors. We are recreating that, but in a way and at a level that works for machines

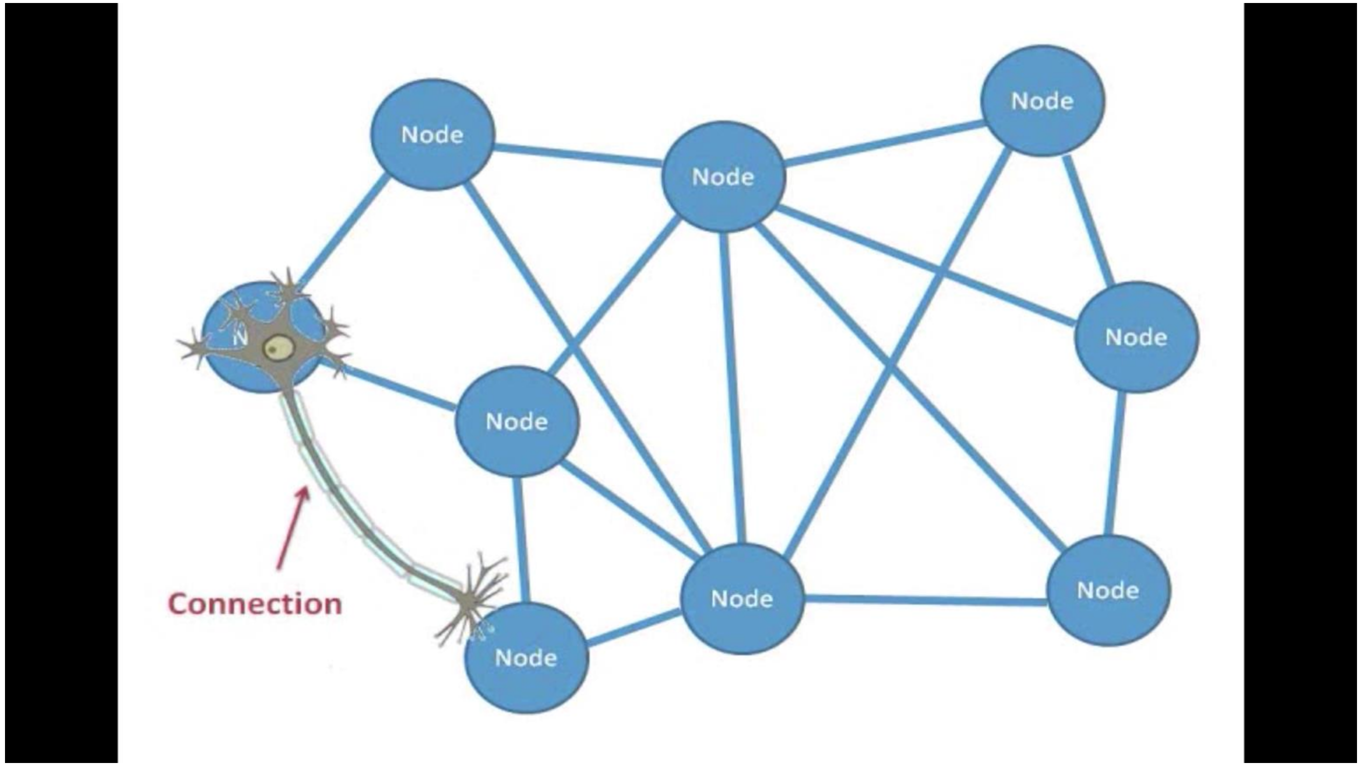

Deep Learning (Neural Networks)

- The idea behind deep learning algorithm, you get input from observation and you put your input into one layer

- That layer creates an output which in turn becomes the input for the next layer, and so on

- This happens over and over until your final output

Deep Learning (Neural Networks)

- Some neural networks can be deeper, based on the number of layers it has!

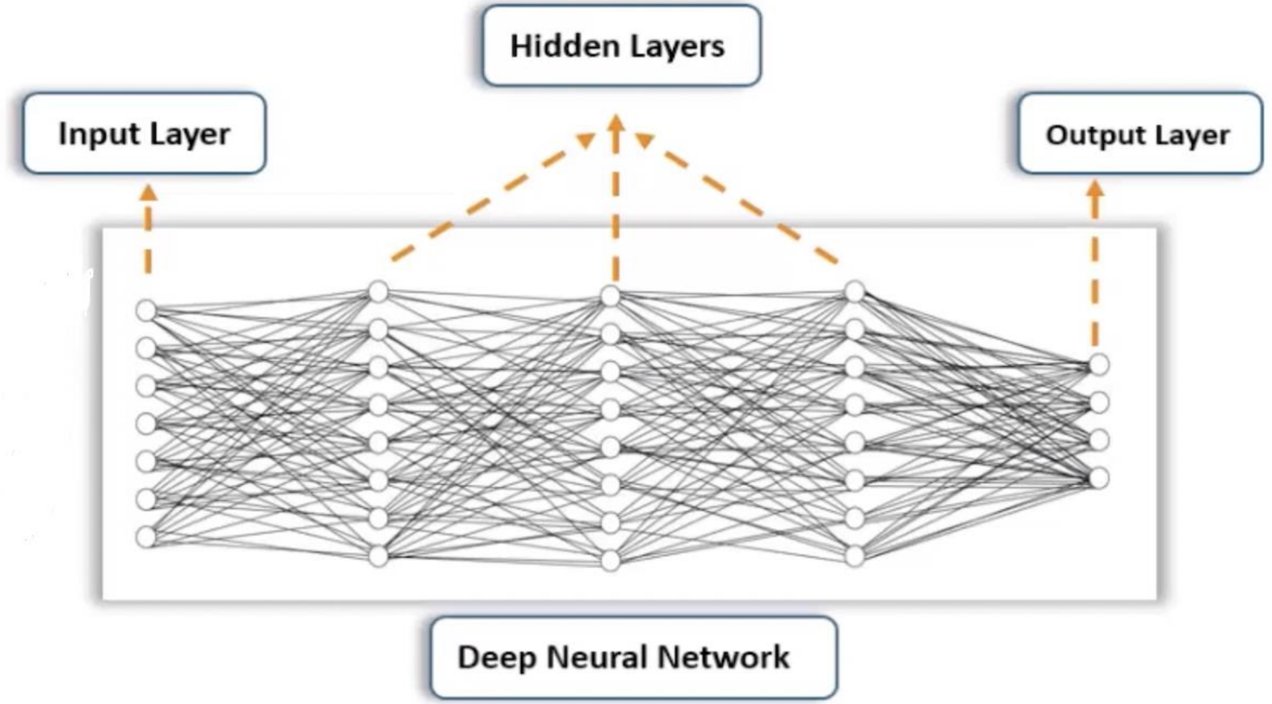

How Neural Networks Work?

- Our end goal is to reduce the output error
- Weight is calculated between the layers and adjusted through back-propagation in every iteration

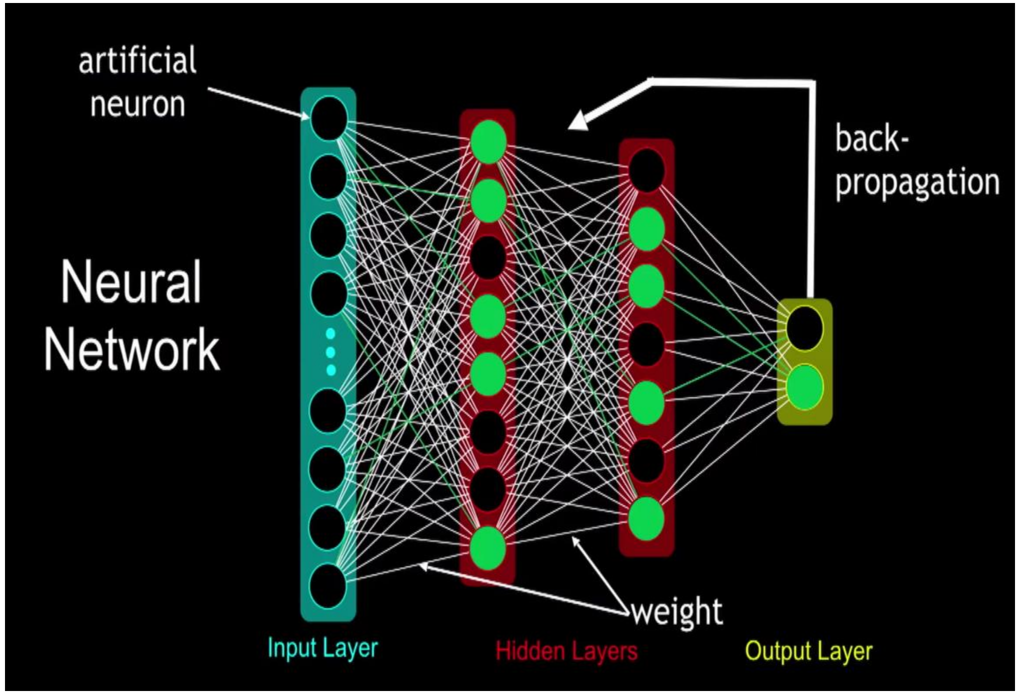

How Neural Networks Work?

- Back-propagation is the essence of neural net training
- It is the practice of fine-tuning the weights of a neural net, based on the error rate obtained in the previous iteration
- Proper tuning of the weights ensures lower error rates

How Neural Networks Work? Example

Let's try to identify circle, rectangle and square

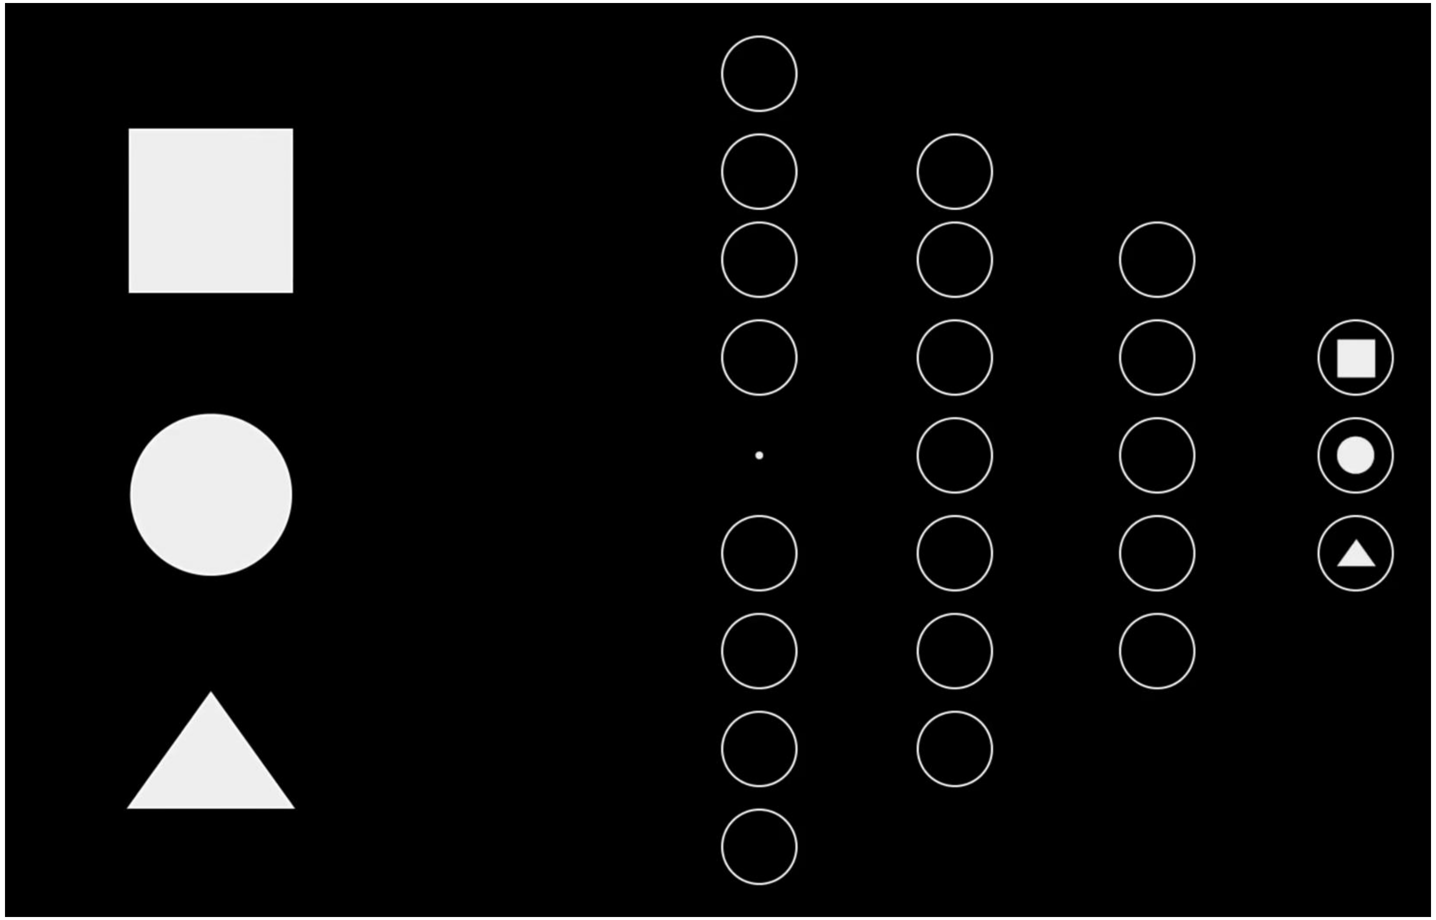

https://t.me/learningnets

How Neural Networks Work? Example

- Let's show circle to our neural network and see whether it can correctly identify it
- Circle is 28x28 = 784 pixels and each pixel is placed as input data to the input layer of our neural network

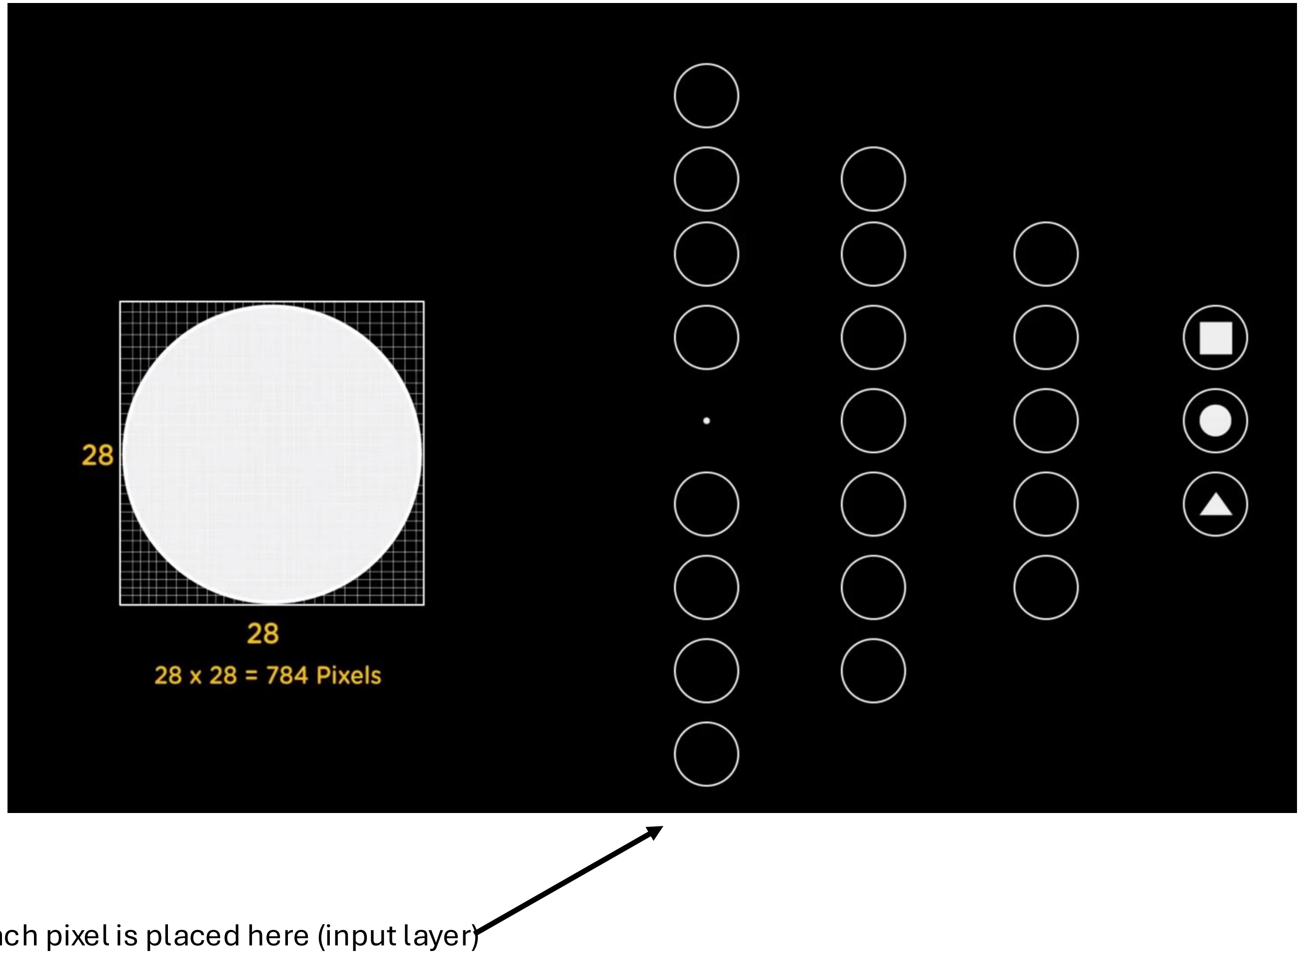

https://t.me/learningnets

How Neural Networks Work? Example

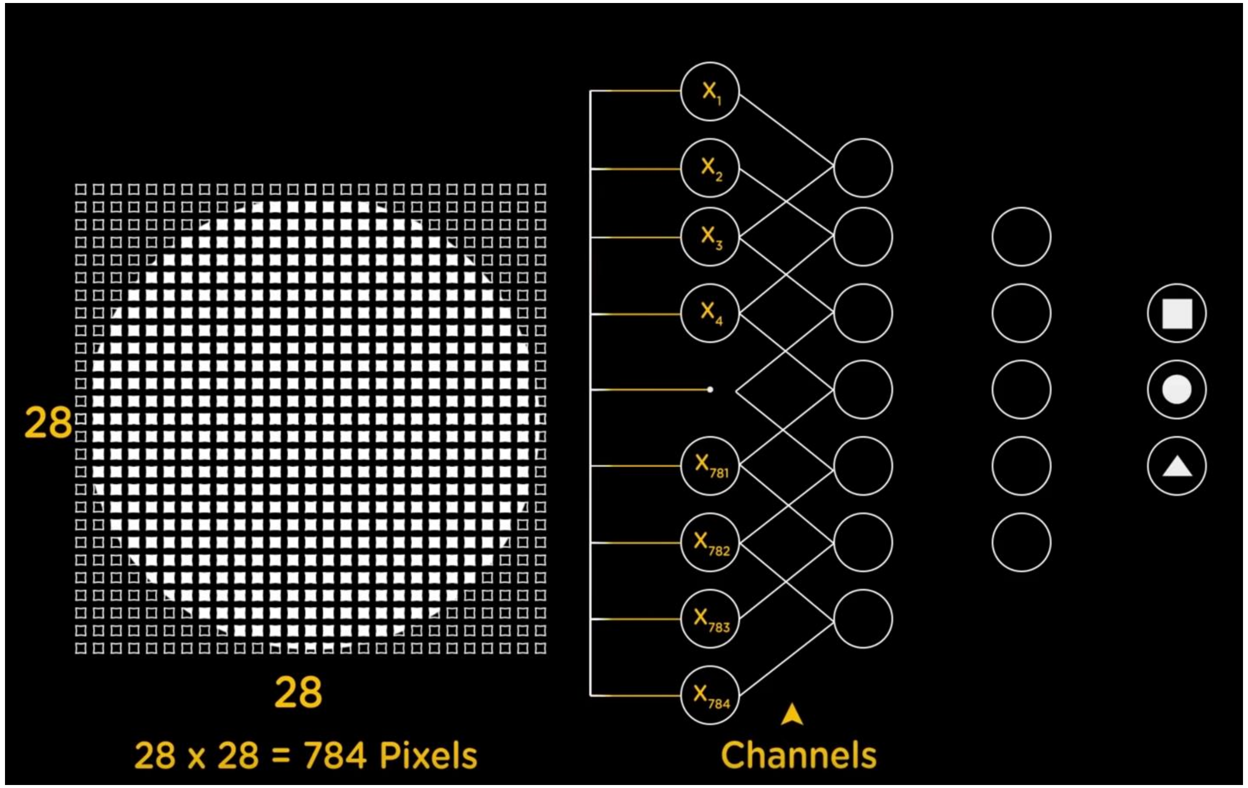

Layers are connected to each other through Channels

How Neural Networks Work? Example

Weight is assigned for each channel initially as random

Both input data and output data is provided to training algorithm during the training process

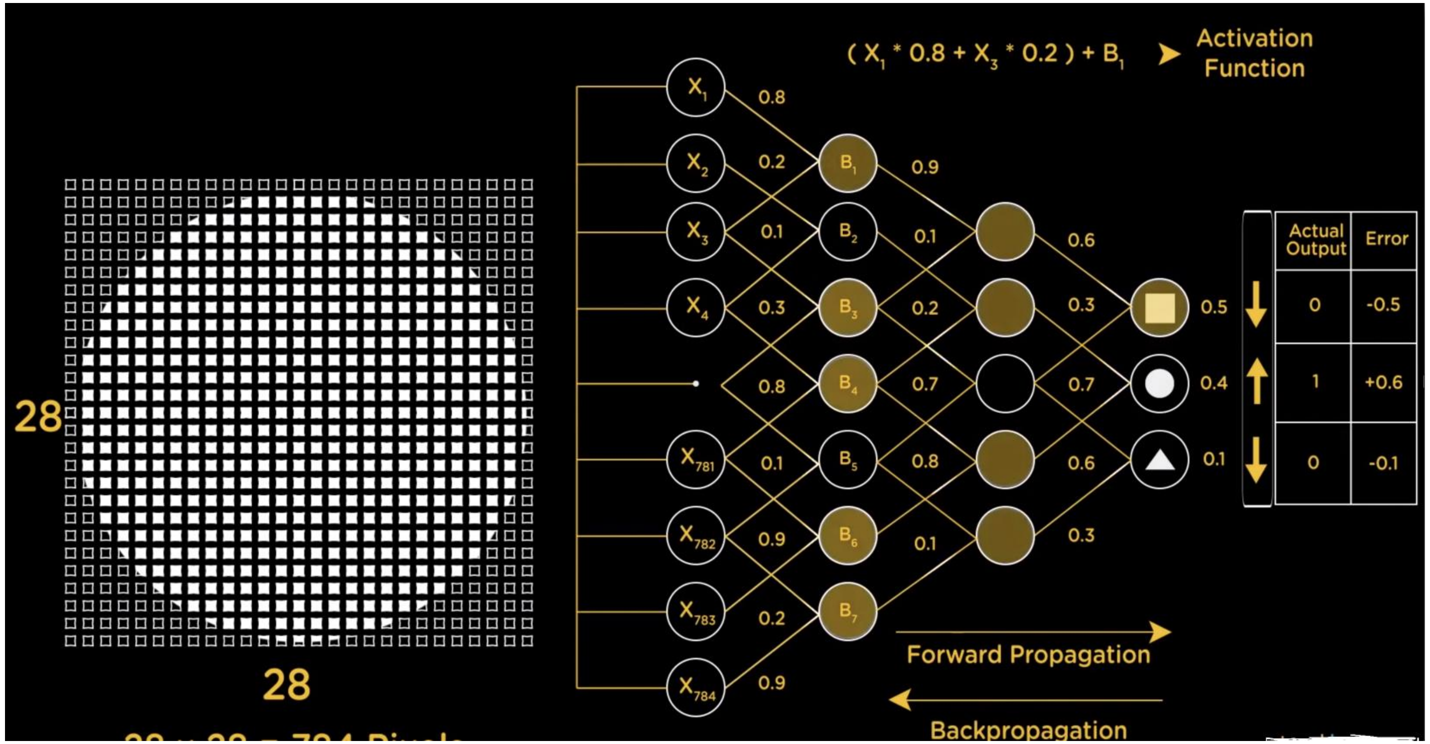

How Neural Networks Work? Example

Activation function is threshold value basically

When the weighted sum exceeds the threshold value, node fires and it will be eligible to send the input to next layer through a forward propagation

When the calculated output – final result is different than actual output, weight is adjusted through back-propagation

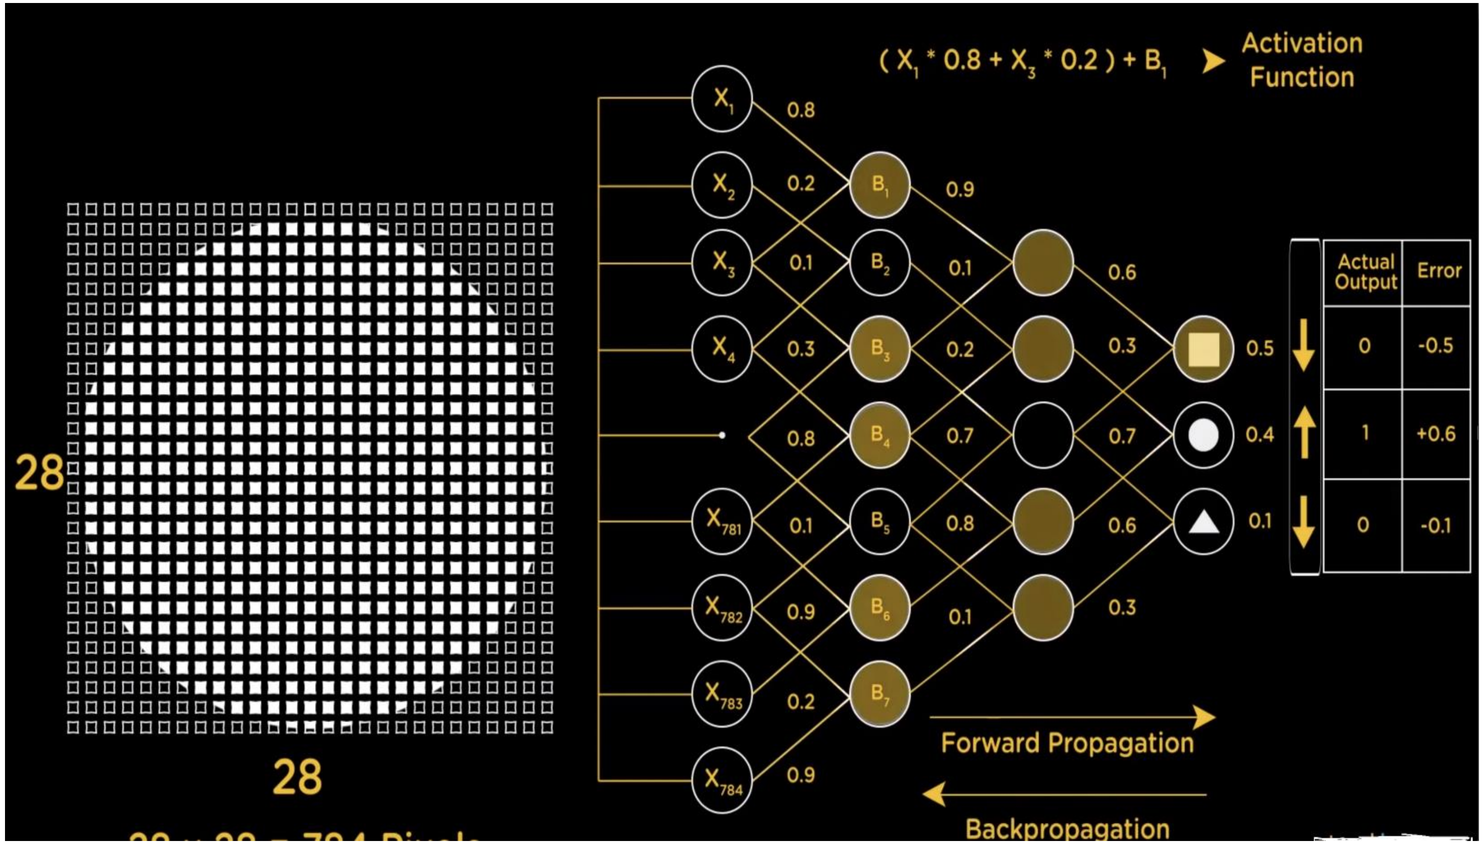

AI/Machine Learning in Network Design

AI/Machine Learning in Network Design

- 3.1 Business Needs
- 3.2 Data Sovereignty
- 3.3 Security
- 3.4 Assurance
- 3.5 Integrity
- 3.6 Impacts
- 3.7 Auto Scalability
- 3.8 Cost and ROI
- 3.9 Governance
- 3.10 Sustainability
- 3.11 AI Network Design Use Cases

Business Needs

For network design, aligning AI with organizational goals means understanding the bigger picture.

It helps predict traffic patterns

Business Needs

AI can analyze historical data and recommend network changes.

Predictive analytics and providing suggestions

AI can recommend right-sizing your resources.

Source: Investopedia

Business Needs

Aligning AI with business needs ensures technology, like intent-based networking, supports organizational goals and constraints effectively.

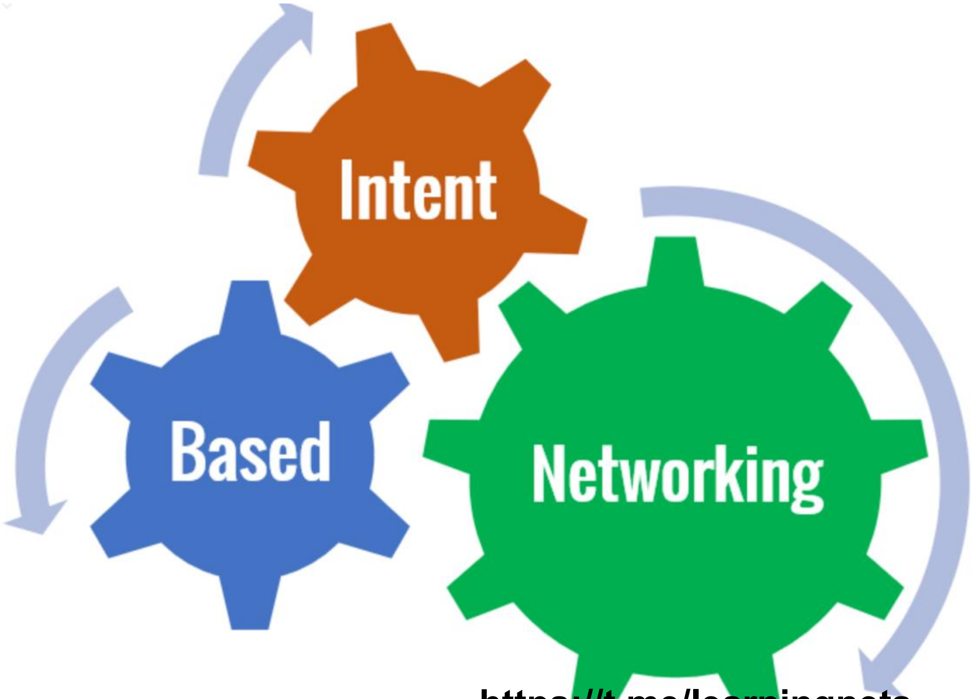

Business Needs

I helped a company with 2000+ branches design their SD-WAN for redundancy, growth, and security.

AI could analyze requirements and constraints to provide accurate, prioritized suggestions

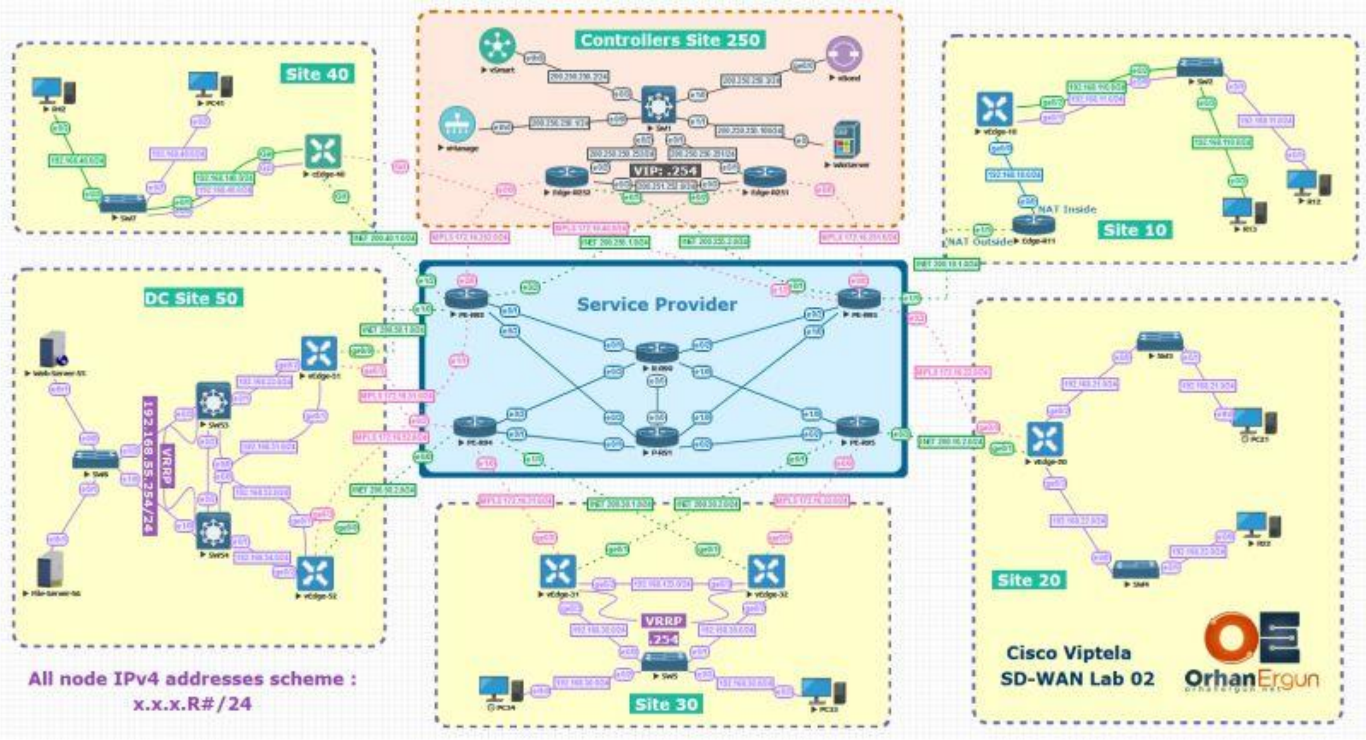

https://t.me/learningnets
Source: OrhanErgun Cisco SD-WAN course

Data Sovereignty

Global businesses rely on the cloud, but local laws apply.

Networks must follow these laws while maintaining performance.

Data Sovereignty

AI can classify and tag data based on sensitivity and compliance requirements.

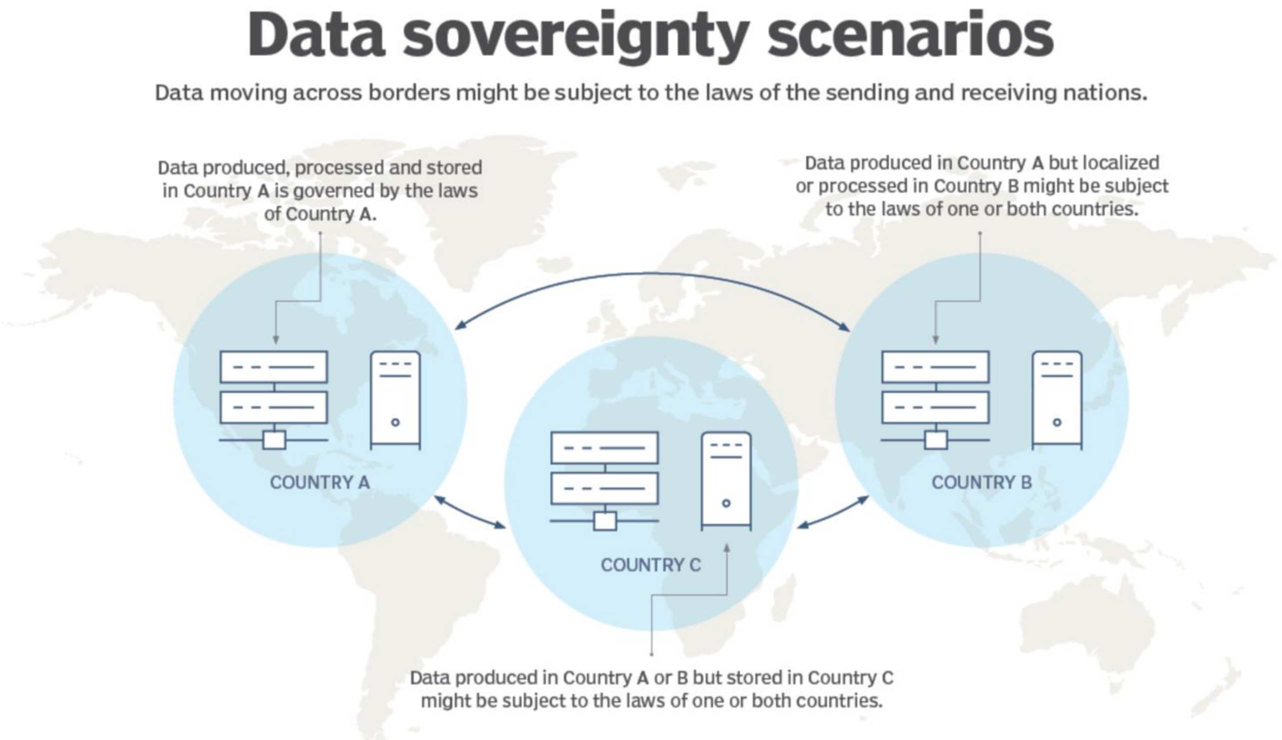

Source: Techtarget

Data Sovereignty

Traffic Engineering is a good example!

AI can ensure that traffic flows comply with regulations: Link Coloring, affinity bits etc.

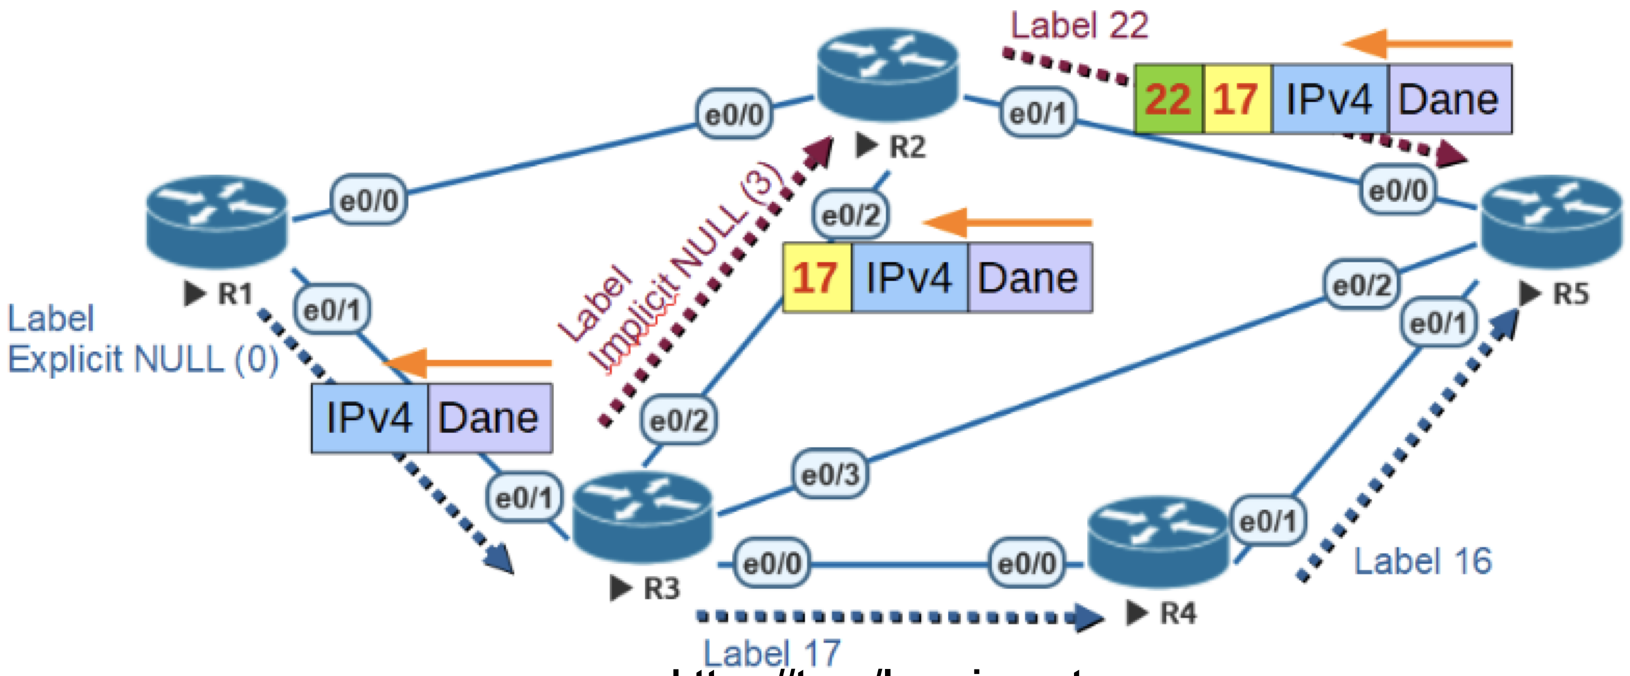

https://t.me/learningnets

Data Sovereignty

AI-powered monitoring

Imagine a scenario where an employee accidentally uploads sensitive European customer data to a U.S. server.

AI can catch that, warn you, or fix it if you give permission.

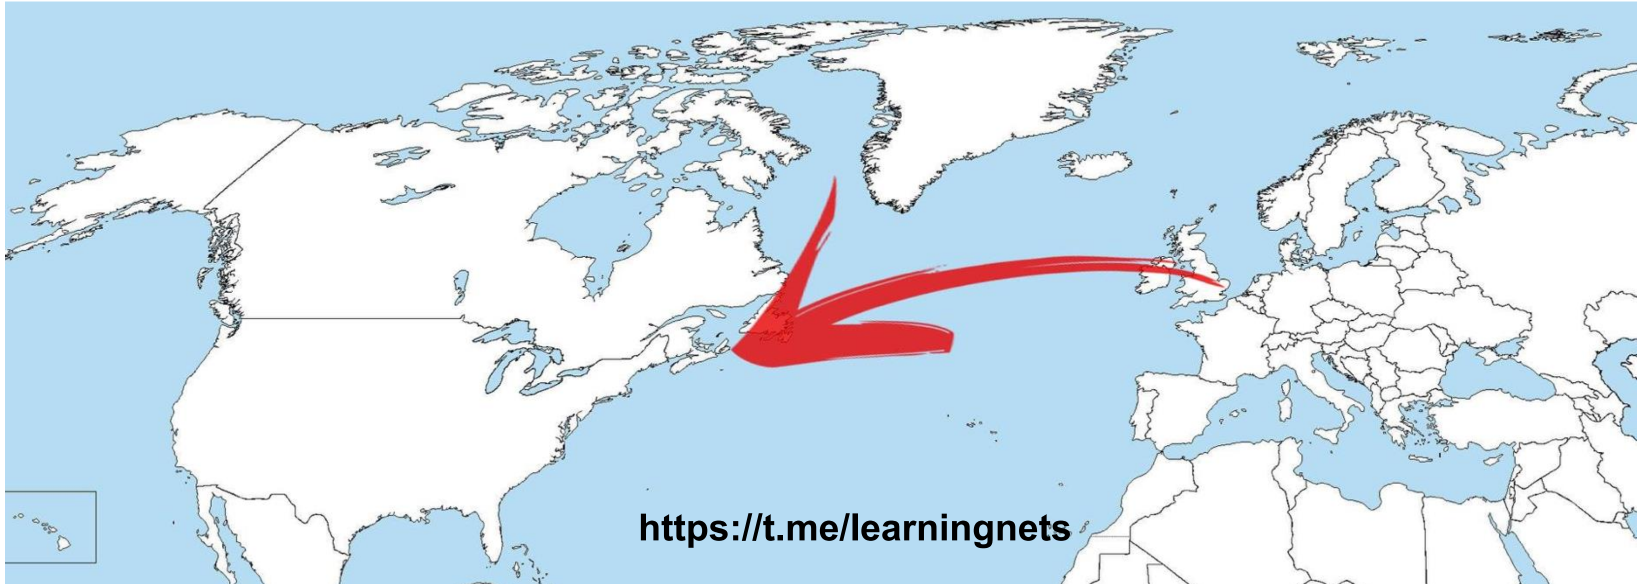

Security

AI is not just making things smarter; it also brings new vulnerabilities.

AI in security, the hottest topic is Anomaly detection

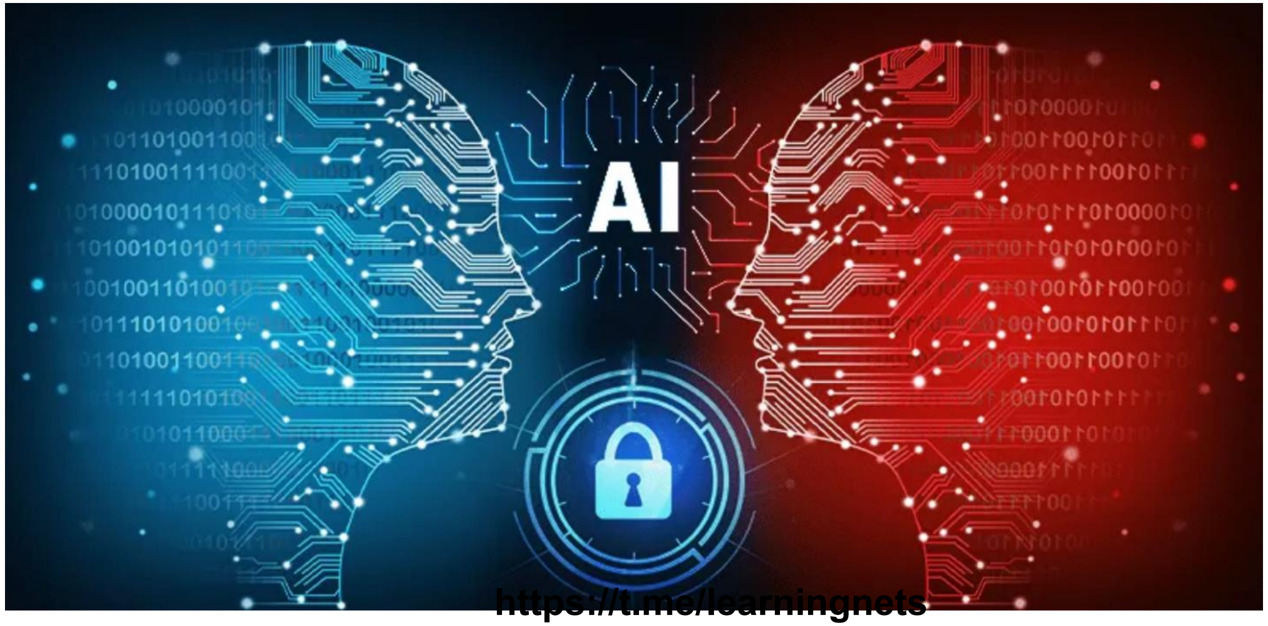

Security

It learns what regular traffic looks like and alerts about anything unusual.

Security

AI can help with incident response too!

For example, if malware infects a device, AI can isolate that device from the network instantly, preventing the spread.

Security

On the other hand, AI systems need to be trained with correct and up-to-date data.

The AI might miss real threats or create false alarms if the data is incomplete or biased.

Assurance

Assurance is guaranteeing the reliability and availability of the networks.

Keeping the network running smoothly in any case.

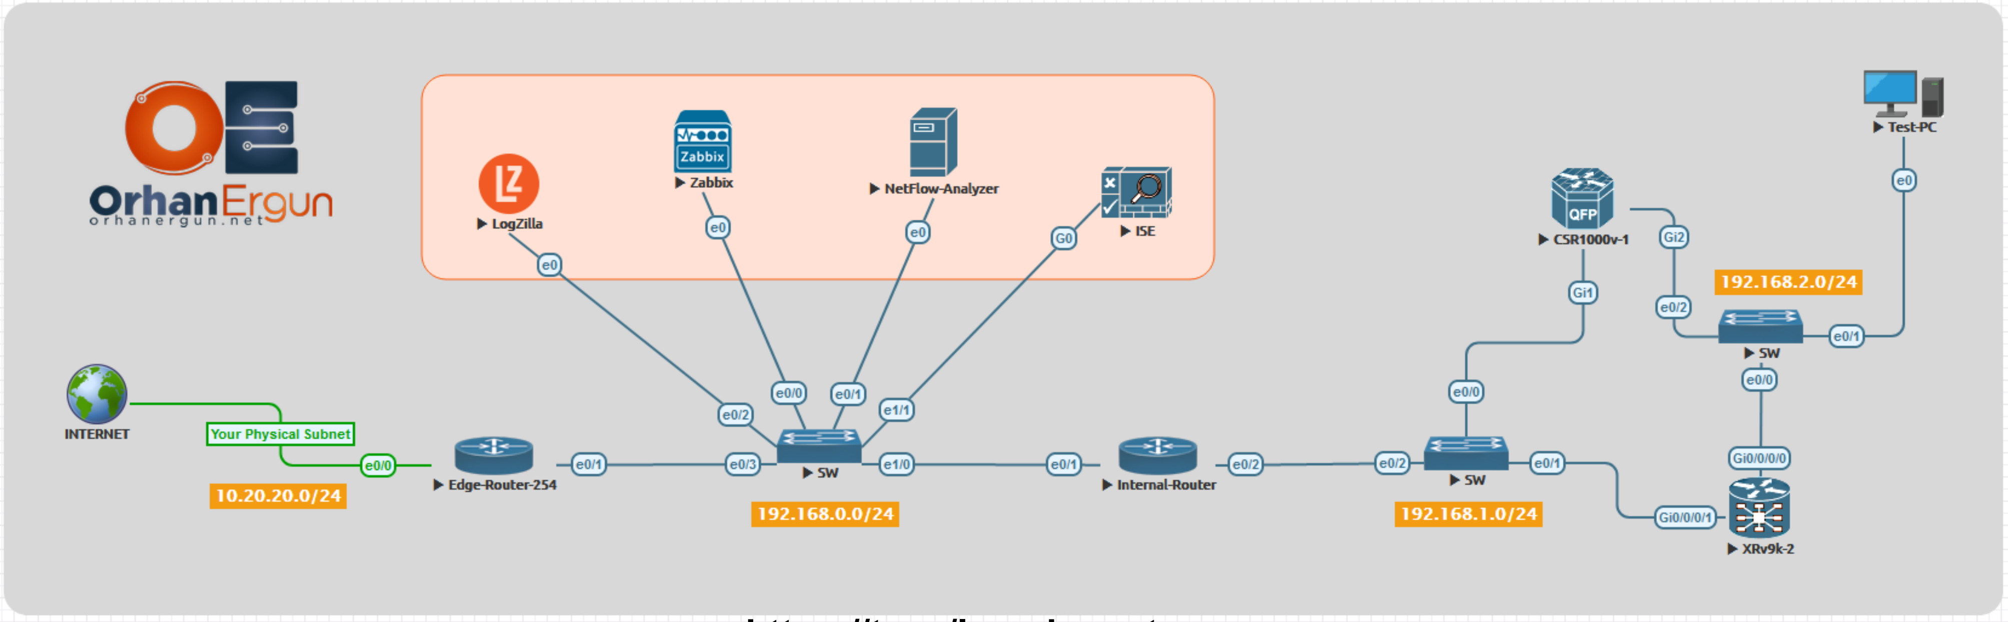

https://t.me/learningnets

Source: Orhan Ergun Network Assurance Course

Assurance

One way to increase networking assurance is by adding redundancy.

But the tradeoff here is cost and complexity, redundancy is not cheap

Assurance

Assurance might prioritize some services and deprioritize others.

Imagine AI dynamically managing traffic for a financial firm.

If the trading traffic gets deprioritized due to an AI miscalculation, it can cost a lot. Think about HFT.

Assurance

With assurance, reliability comes, but cost and complexity increase due to redundancy and multiple components

Integrity

Integrity means Ensuring Data Accuracy and Consistency.

Suppose your AI is recommending network changes for peak-hour traffic management.

If the input data is corrupted, let’s say, it shows outdated traffic patterns, it might suggest routing traffic through already congested links.

Integrity

This is why data validation is very important.

AI could recommend moving sensitive patient data to a cheaper cloud storage option.

If the AI doesn’t understand compliance rules, like HIPAA, this could lead to legal violations.

Integrity

If AI wrongly flags legitimate traffic as a threat, it could block critical business operations.

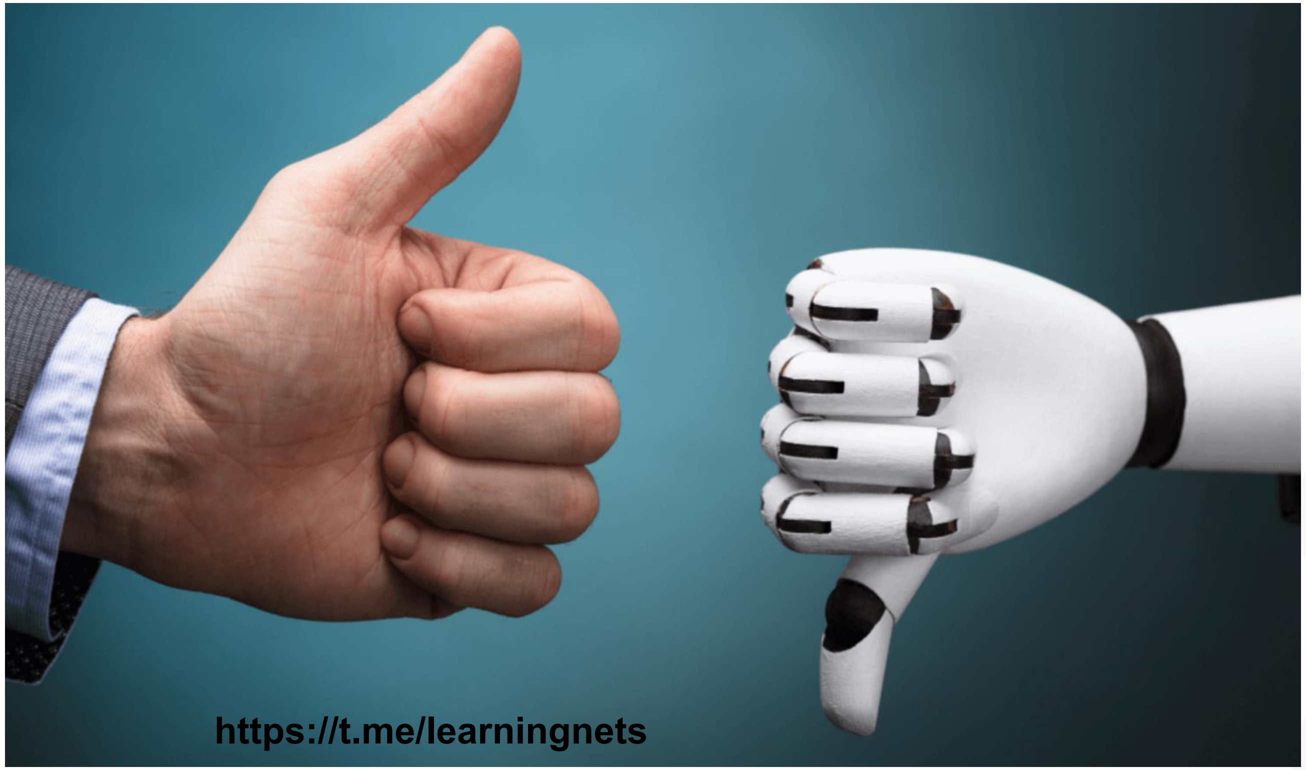

Integrity

Tradeoff:

Validation and data checking take time, but:

They prevent bad decisions from affecting the network

Ensure compliance with rules.

Build trust in AI systems!

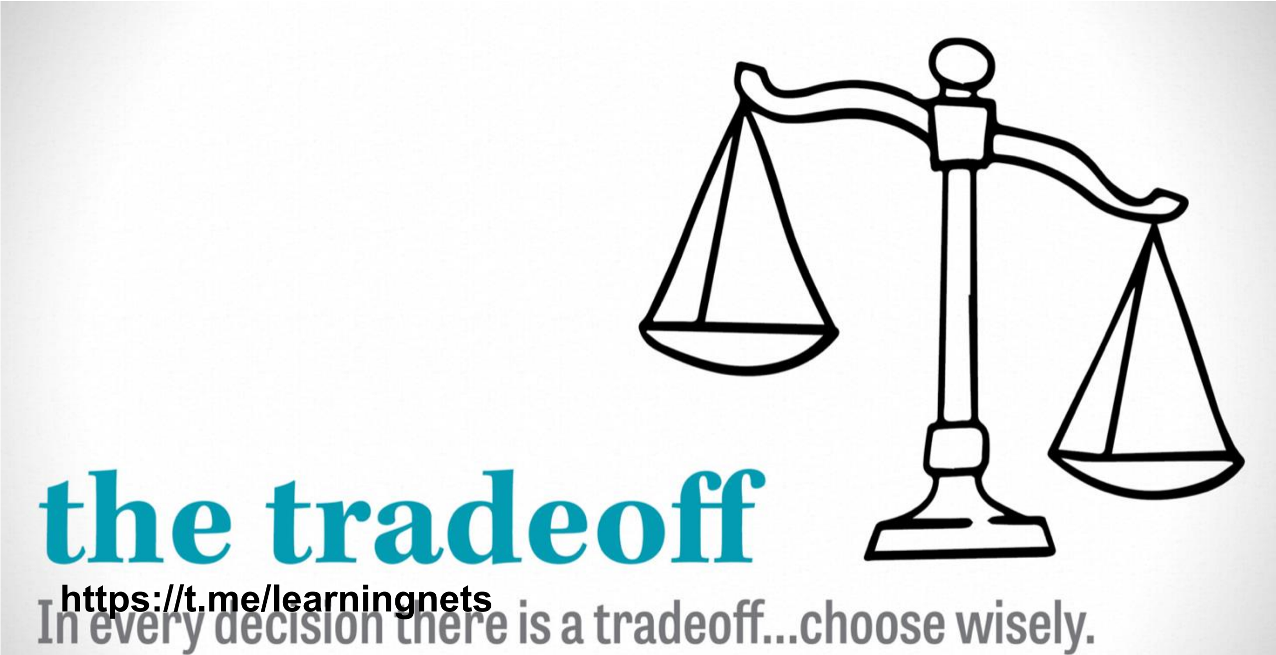

Impact

Storage and traffic pattern impact:

Telemetry data, network logs, traffic analytics and of course, all have to be stored somewhere.

It is used for training models, creating baselines, or even troubleshooting issues.

Impact

You may redesign storage by adding high-performance systems or offloading less critical data to the cloud.

AI changes traffic patterns; for IoT setups with constant data streaming, use Edge AI to cut bandwidth needs to central locations.

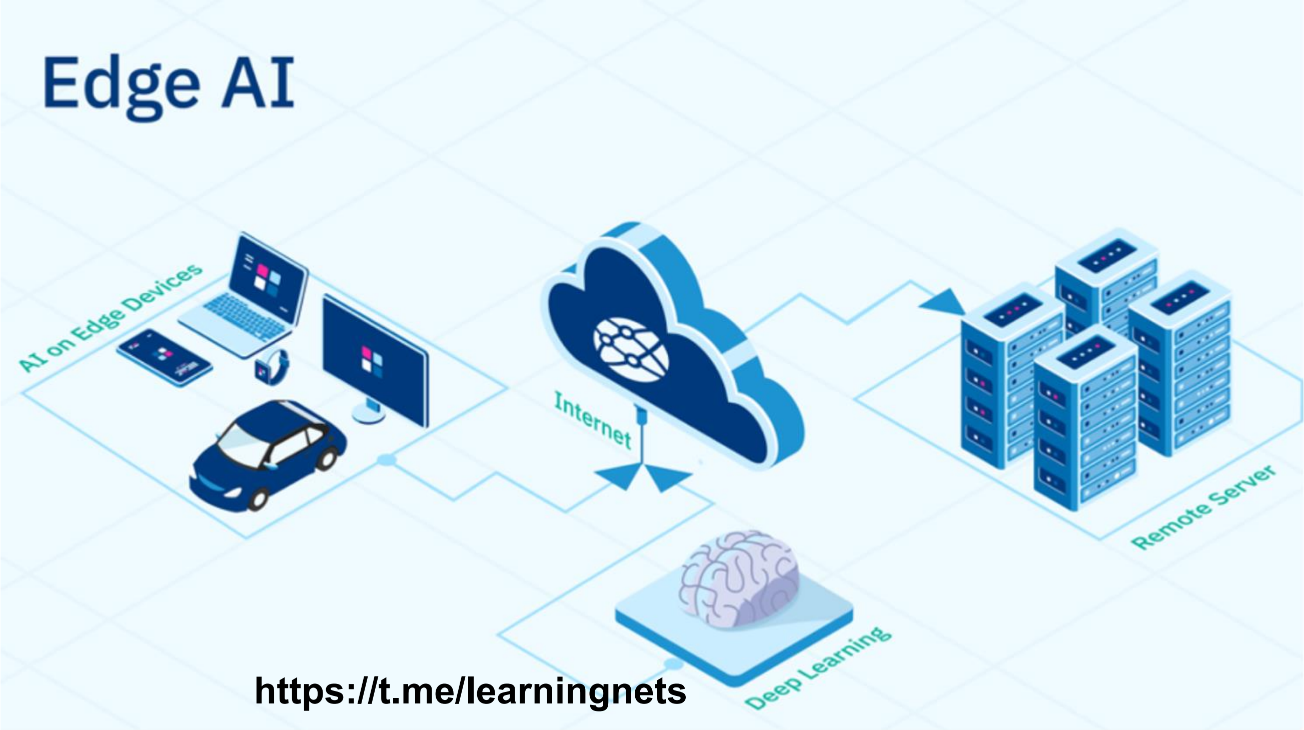

Impact

Summary

With more data, more insight, and better network performance, the drawbacks are more storage requirements, more bandwidth, etc.

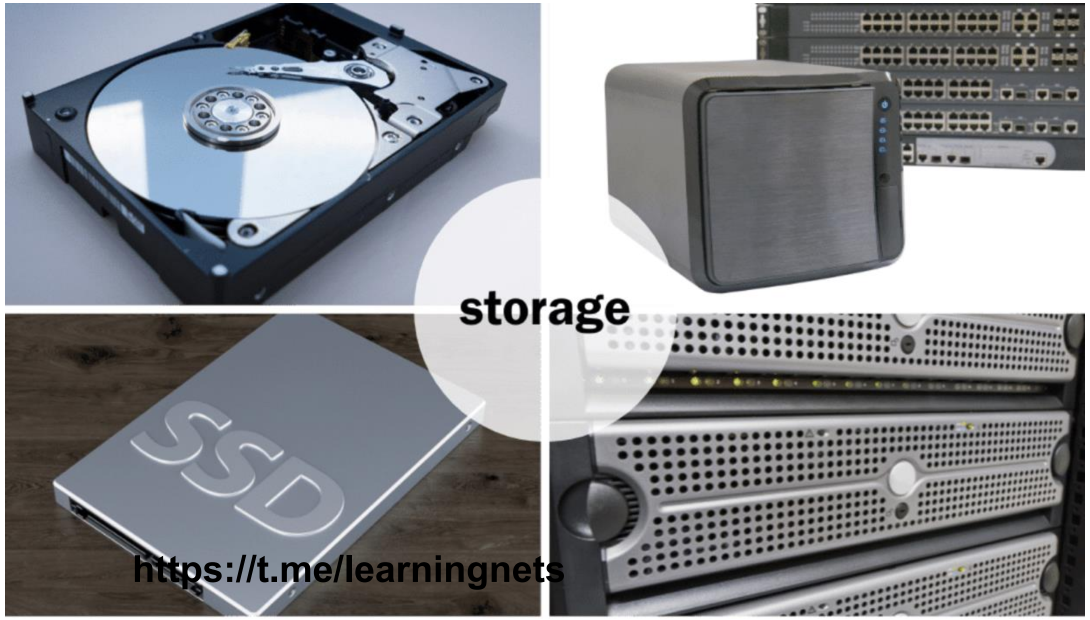

Auto-Scalability

Imagine a network that scales itself dynamically based on current demand.

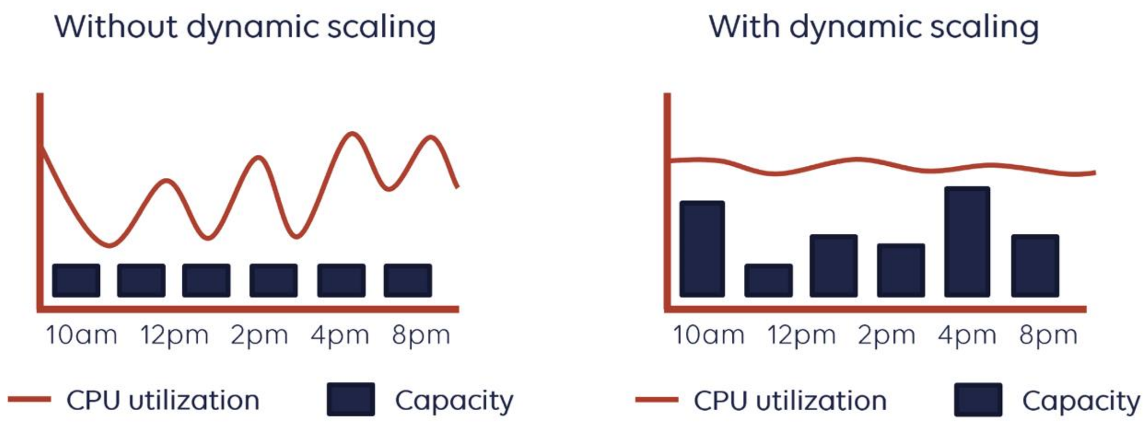

Source: stormit

Auto-Scalability

On Black Friday, your e-commerce network sees a 300% traffic surge.

AI-driven scaling quickly adds resources like VMs, bandwidth, or compute power where needed.

Auto-Scalability

AI analyzes historical data and trends to enable predictive auto-scaling.

While auto-scaling is an impressive feature, it requires extensive backend preparation.

The benefits include enhanced performance, optimized resource usage, and, most importantly, satisfied users

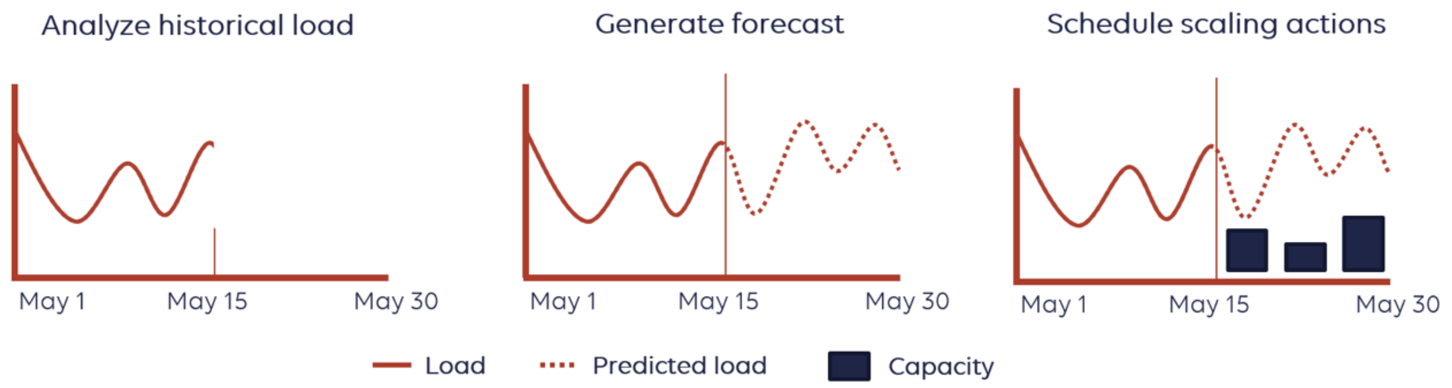

Source: stormit

Cost and ROI

Integrating AI into your network comes with significant costs.

Upfront investment, or CAPEX, includes purchasing AI tools, upgrading infrastructure, and training your team.

Cost and ROI

AI has ongoing costs for system maintenance and data storage but can save money long-term if implemented properly.

Cost and ROI

AI automates tasks like configuration management and fault detection, allowing network engineers to focus on strategic initiatives

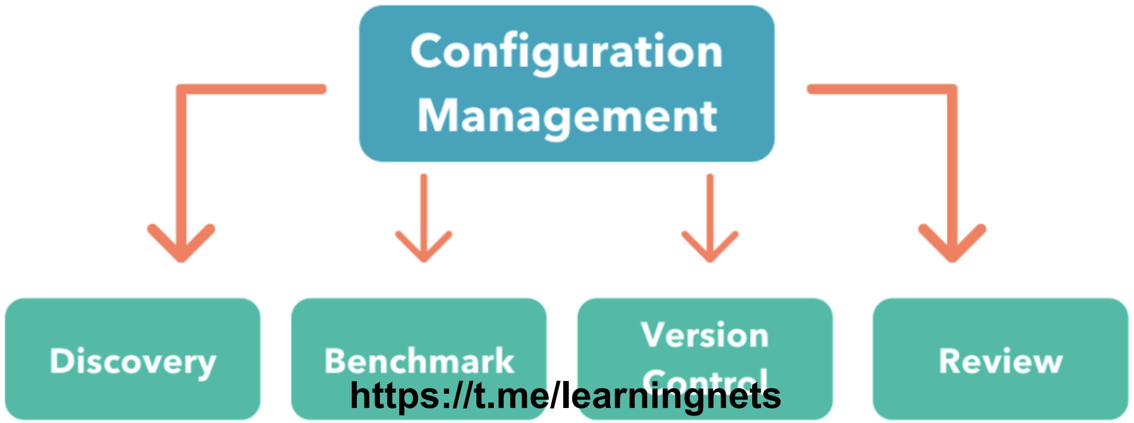

Cost and ROI

From an ROI perspective, faster troubleshooting reduces downtime, boosts revenue, and improves customer satisfaction.

AI's predictions help prevent costly outages, benefiting industries like banking, e-commerce, and beyond.

Governance

Governance in networking involves creating policies for AI use, covering ethics, operations, and regulatory compliance.

Source: boardeffect

Governance

For instance, an AI system managing traffic could mistakenly deprioritize critical applications without proper policies in place.

Governance

Transparency is essential in Governance.

For example, if your AI suggests a network change, it should explain why.

Explainable AI, XAI!

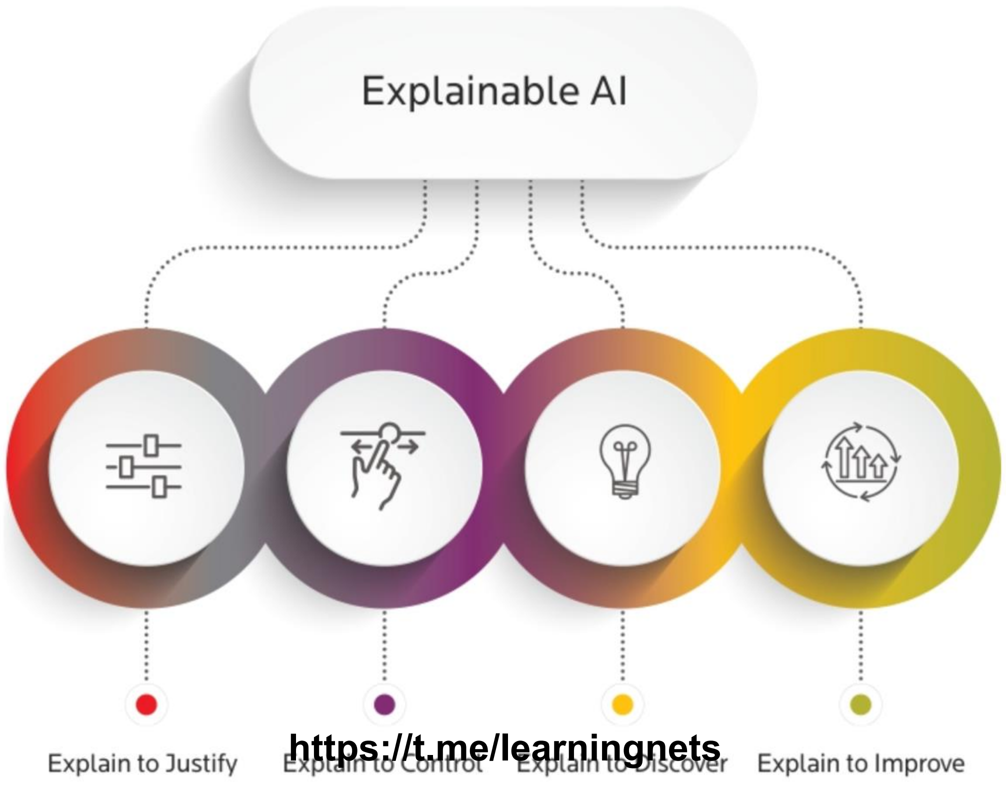

Governance

The drawback is It adds complexity and can slow down decision-making.

Sustainability - Green AI

An essential aspect of AI deployment is sustainability, which focuses on minimizing the environmental impact of building and running AI systems.

High computational demands translate to significant energy consumption.

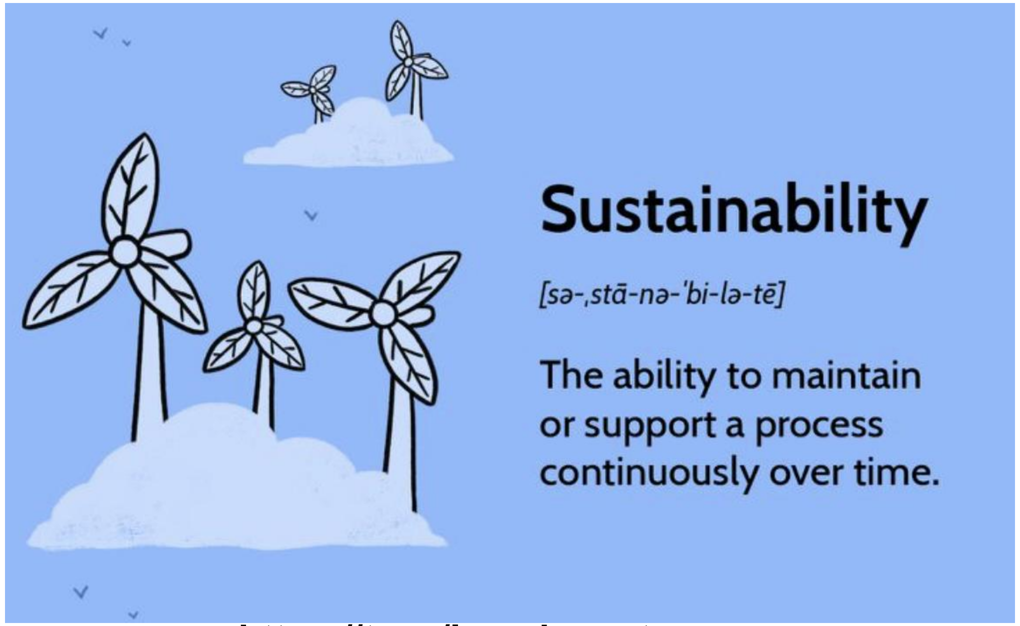

Source: Investopedia

Sustainability - Green AI

Data centers, the backbone of networking, use vast amounts of electricity for both servers and cooling.

Training LLM models, requiring tens of thousands of GPUs, demands so much power, including for cooling systems.

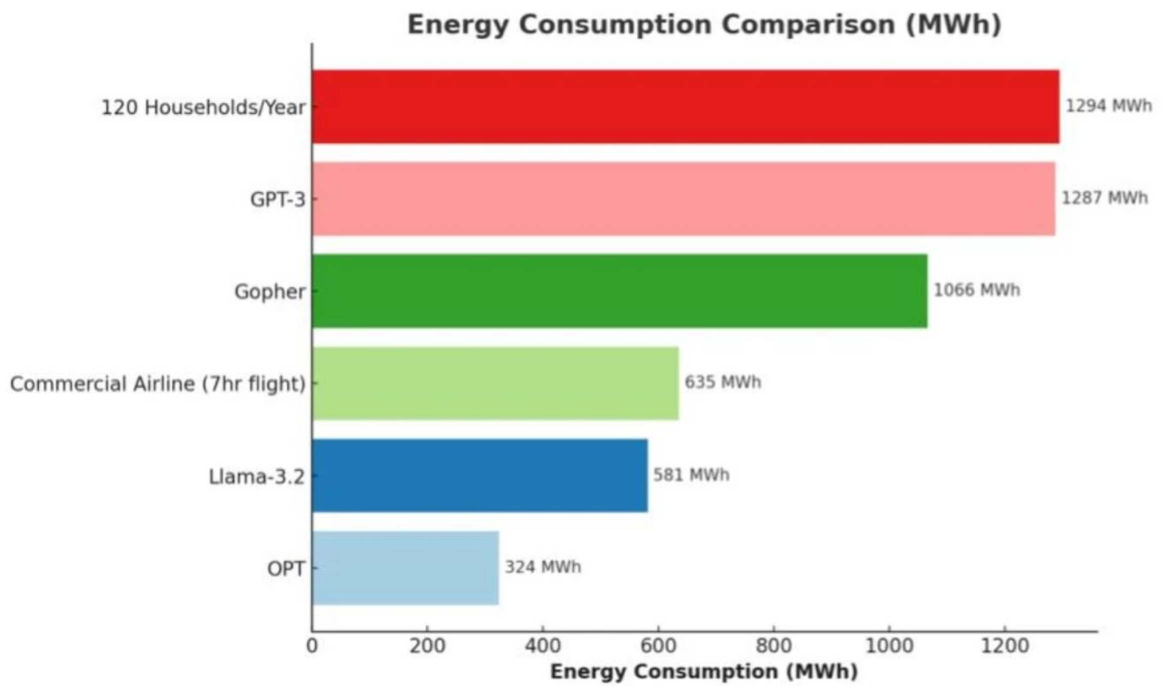

Source: computer.org

Sustainability - Green AI

Google's DeepMind uses AI to control cooling systems in their DCs, reducing energy consumption by 30%.

Sustainability - Green AI

Green AI is all about designing AI models and processes that are specifically optimized for energy efficiency.

Because training a large AI model can have a huge carbon print

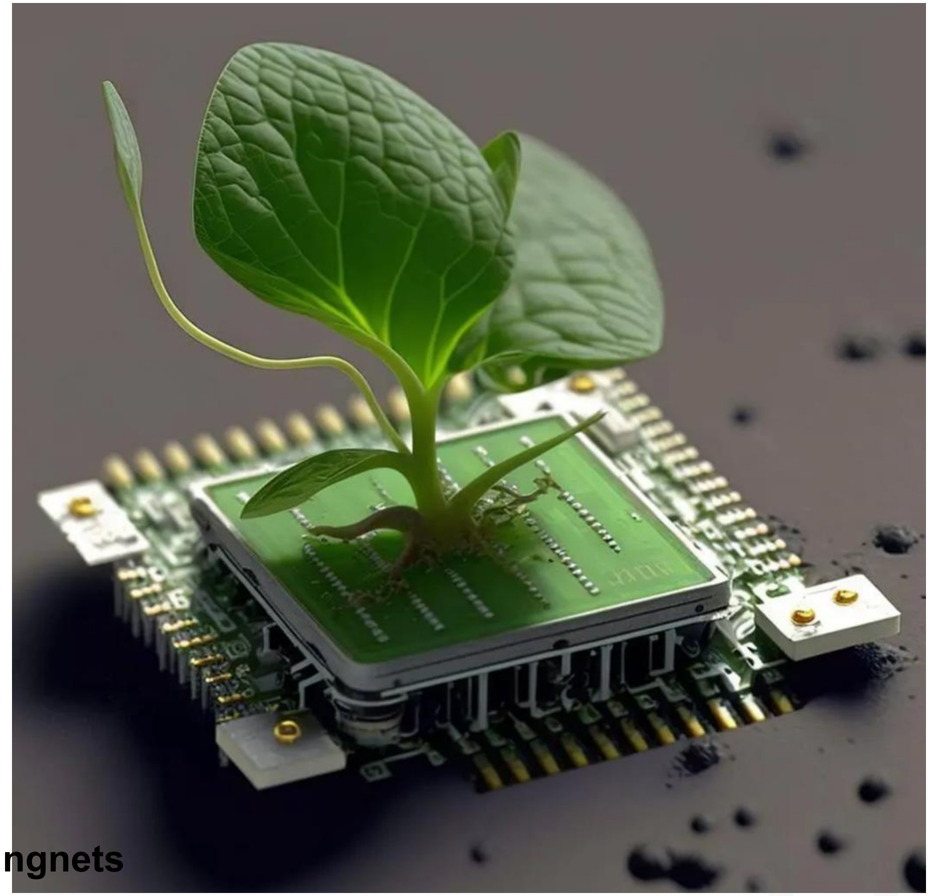

Sustainability - Green AI

For an AI-based traffic engineering system, instead of using a large, complex model with terabytes of data, you can fine-tune smaller, efficient models to achieve similar results with less energy.

This is called Fine-Tuning.

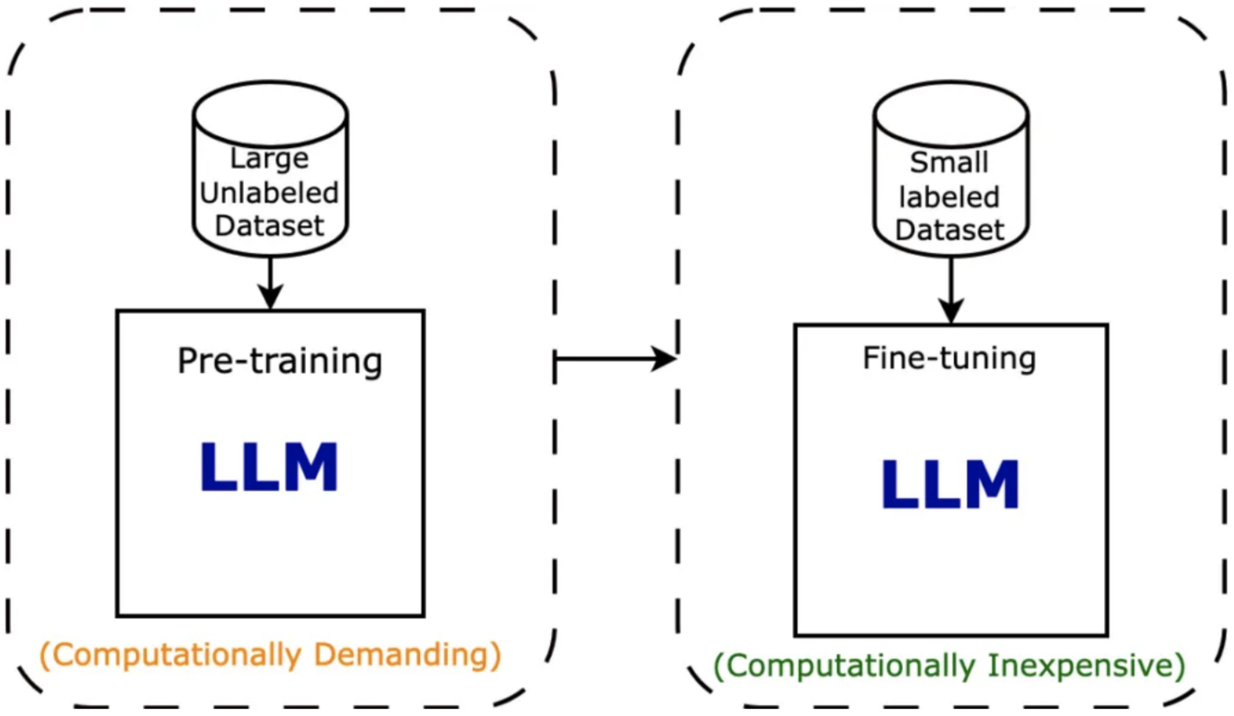

Source: intuitive tutorials

Sustainability - Green AI

It is all about building AI systems that are efficient and sustainable.

Can we do the tasks with less power and resources?

Sustainability – Affordability

AI is not cheap!

It needs hardware like GPUs, TPUs, or other accelerators to perform its heavy computations.

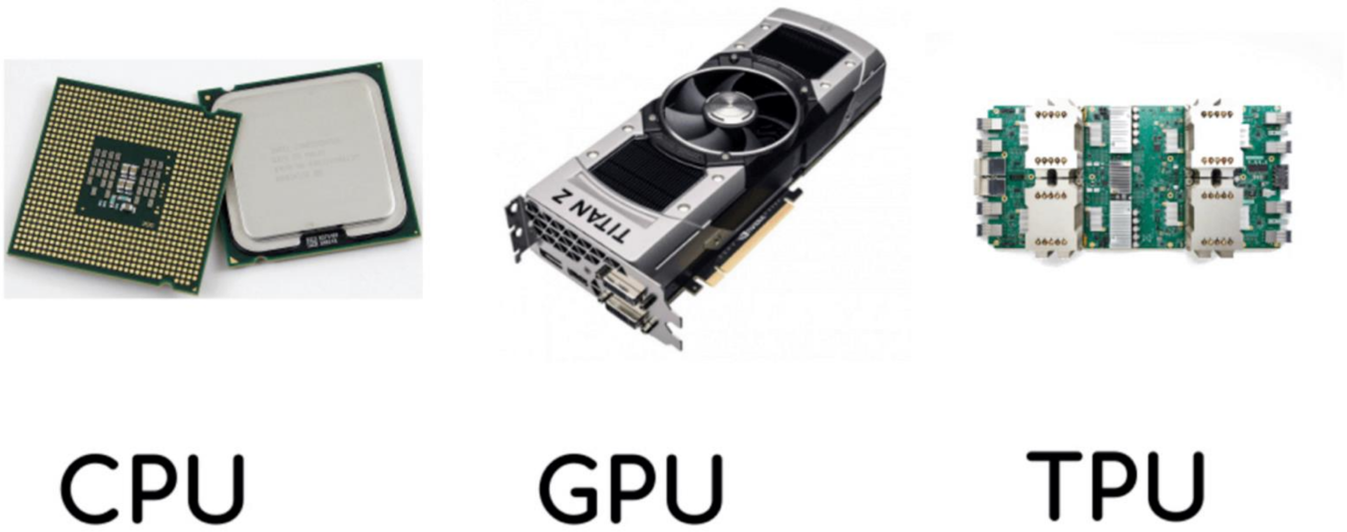

Sustainability – Affordability

CPU/GPU tweaking is using the right hardware for the right use case

Offload non-critical AI tasks to CPUs or energy-efficient AI chips instead of relying solely on GPUs for all workloads.

Sustainability – Affordability

Tweaking hardware by using CPUs for lighter AI tasks and reserving GPUs for demanding processes like real-time anomaly detection.

The balance here is between cost and performance, as GPUs deliver faster results but are more expensive

.

Sustainability – Effective use of Accelerators

Accelerators like TPUs and FPGAs are specialized chips built for AI, offering faster, more efficient training and inference than CPUs.

.

Sustainability – Effective use of Accelerators

These are often cheaper and more energy-efficient than GPUs.

For example, NVIDIA uses GPUs and DPUs in AI-powered network solutions.

DPUs, used in Edge AI, optimize data flows in HPC networks.

.

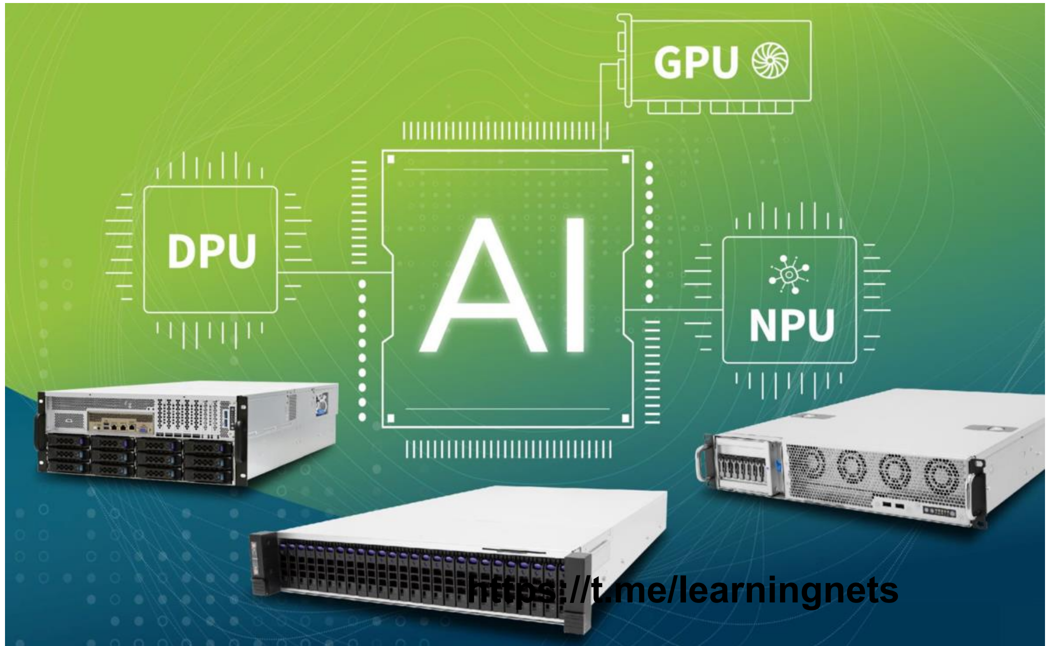

Sustainability – Power and Cooling Requirements

With growing networks and AI, power and cooling are key challenges.

GPUs generate heat, risking overheating without proper design.

Data centers must invest in advanced cooling.

Sustainability – Power and Cooling Requirements

Liquid cooling is becoming popular with massive GPU clusters but is still emerging. Cooling itself also consumes energy.

Boiling liquid carries away heat generated by computer servers at a Microsoft datacenter. Microsoft is the first cloud provider to run two-phase immersion cooling in a production environment. Photo by Gene Twedt for Microsoft.

Sustainability – Power and Cooling Requirements

Advanced cooling systems and power management tools can be expensive to deploy.

And while AI helps optimize resource usage, the upfront costs can be a barrier for smaller organizations.

Sustainability – Power and Cooling Requirements

Your network devices, switches, routers, and cables are key parts of the picture.

High-speed connections like 400G or even 800G are good, but they also consume so much power and generate significant heat.

Sustainability – Power and Cooling Requirements

Another thing to think about is power distribution.

AI-driven systems often have unpredictable power demands.

For example, imagine your AI models suddenly analyzing traffic spikes during peak hours.

Your network devices handling this data might draw more power during these times.

Sustainability – Power and Cooling Requirements

Design must manage power spikes, ensuring systems aren't overloaded or wasting energy.

Edge computing can offload tasks by processing data closer to its source instead of relying entirely on central data centers. (Edge AI)

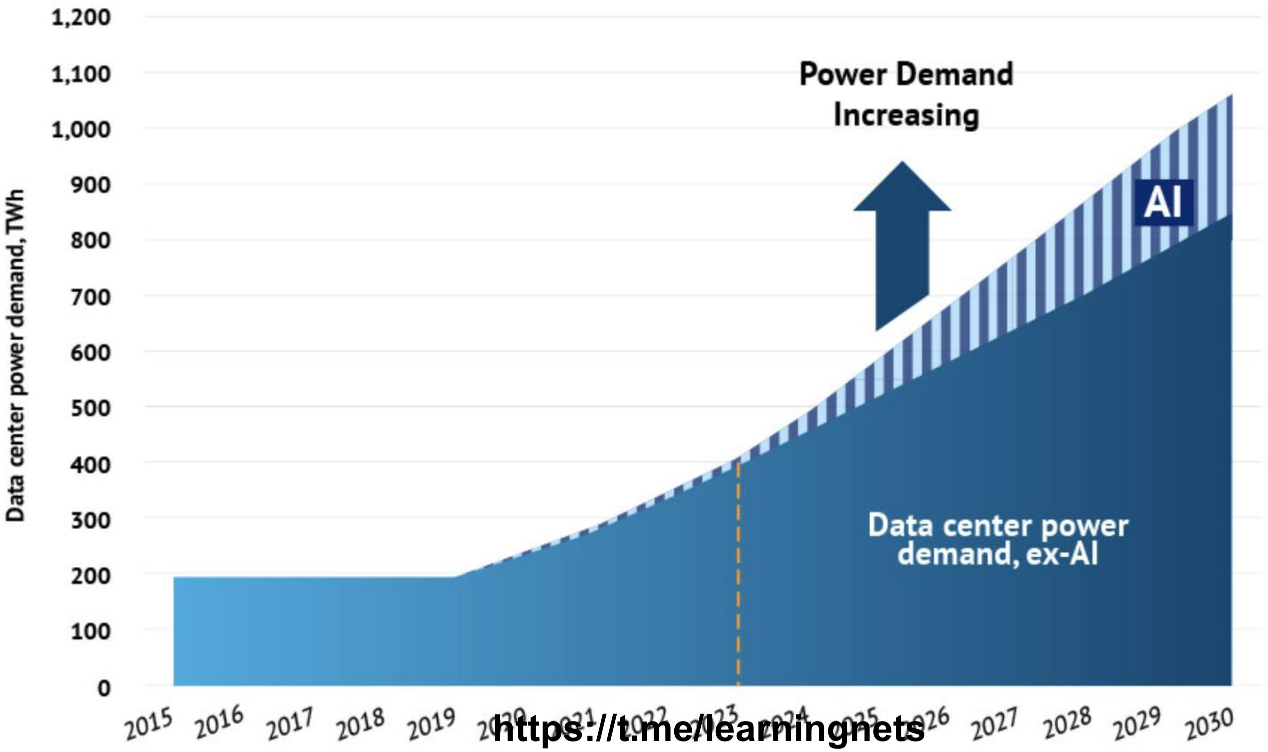

AI Network Design Use Cases- Learning for Predictive Network Modeling

Machine learning enhances predictive network modeling by analyzing data, finding patterns, and forecasting behavior, preventing issues before they occur to minimize downtime and costs.

.

Source: Investopedia

AI Network Design Use Cases- Learning for Predictive Network Modeling

For example, machine learning algorithms can analyze network traffic data to predict peak network traffic usage times, ensuring resources are allocated efficiently to handle increased demand.

AI Network Design Use Cases- Learning for Predictive Network Modeling

Fault detection in ML is used to find anomalies like problematic links or devices by constantly monitoring logs, packet flows, and metrics.

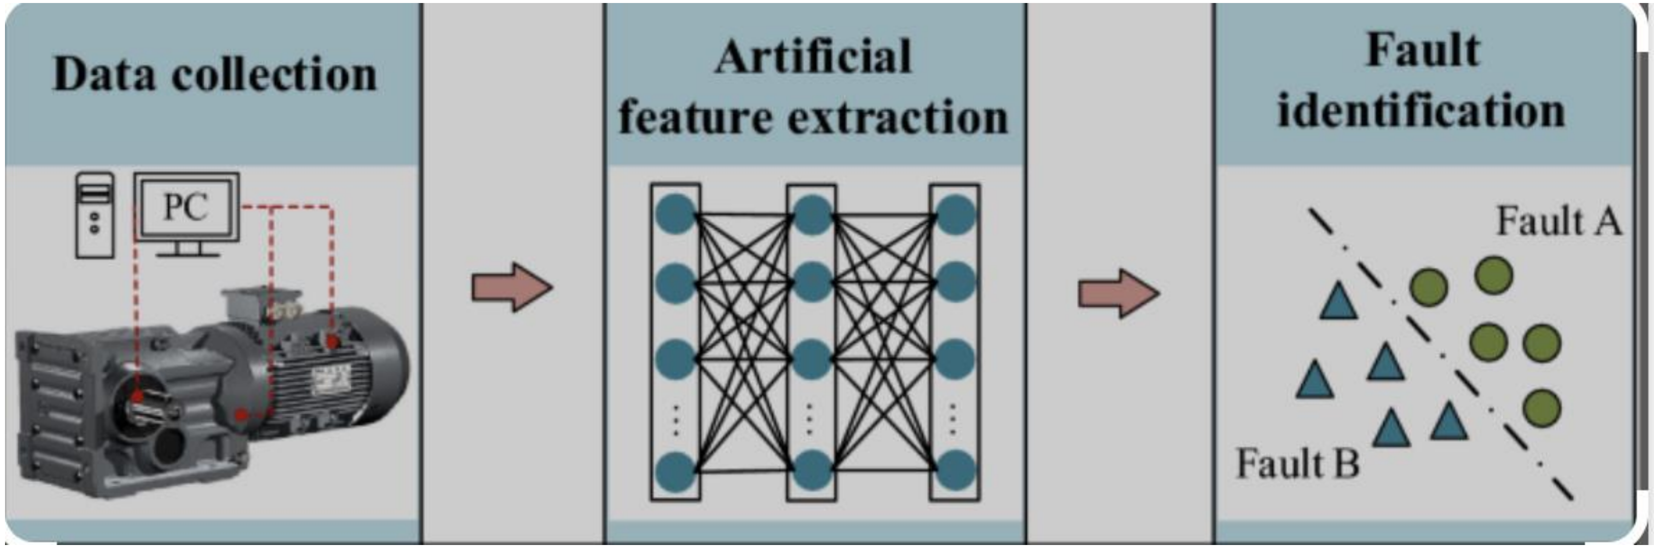

Source: Springer

AI Network Design Use Cases- Learning for Predictive Network Modeling

Time-series analysis is a common technique in predictive network modeling.

This machine learning model uses historical performance trends to predict future metrics like latency, jitter, or packet loss.

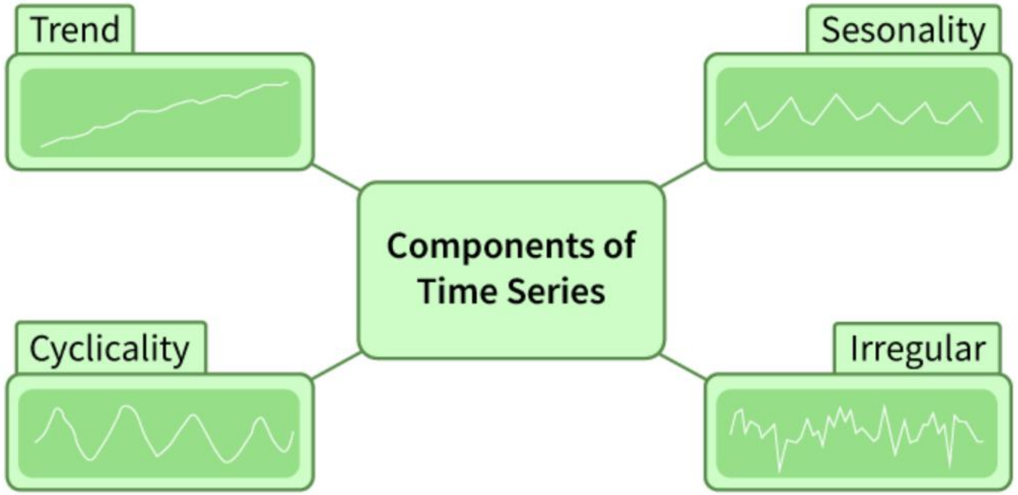

Source: geeksforgeeks.org

AI Network Design Use Cases- Learning for Predictive Network Modeling

Another important use is capacity planning.

Machine learning helps predict when networks will hit resource limits based on growth trends, guiding admins to plan upgrades

.

AI Network Design Use Cases- Learning for Predictive Network Modeling

Predictive modeling enhances traffic engineering by recommending routes based on congestion estimates, using historical trends to suggest alternate paths.

AI Network Design Use Cases- Learning for Predictive Network Modeling

With ML, what-if analysis is easy!

With Predictive network modeling, we can react to problems proactively in ML!

.

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

LLMs like GPT, BERT, and others are revolutionizing network design and automation by leveraging their ability to process, analyze, and generate text-based insights.

.

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

In network design, LLMs can analyze massive volumes of preparatory documents, config guides, validated designs and suggest something.

-

Source: Clockwise

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

Automation is another area for LLMs.

Specify high-level intent, and it can translate into technical configurations for routers, switches, or firewalls.

.

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

If you need redundancy in MPLS network config, LLM can immediately provide it within seconds based on its vendor.

More than config, LLMs can correlate error logs, performance metrics, and historical data to identify the root cause of issues.

.

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

LLMs can help with training and skill development.

They can create training materials, quizzes etc. , for specific network environments.

.

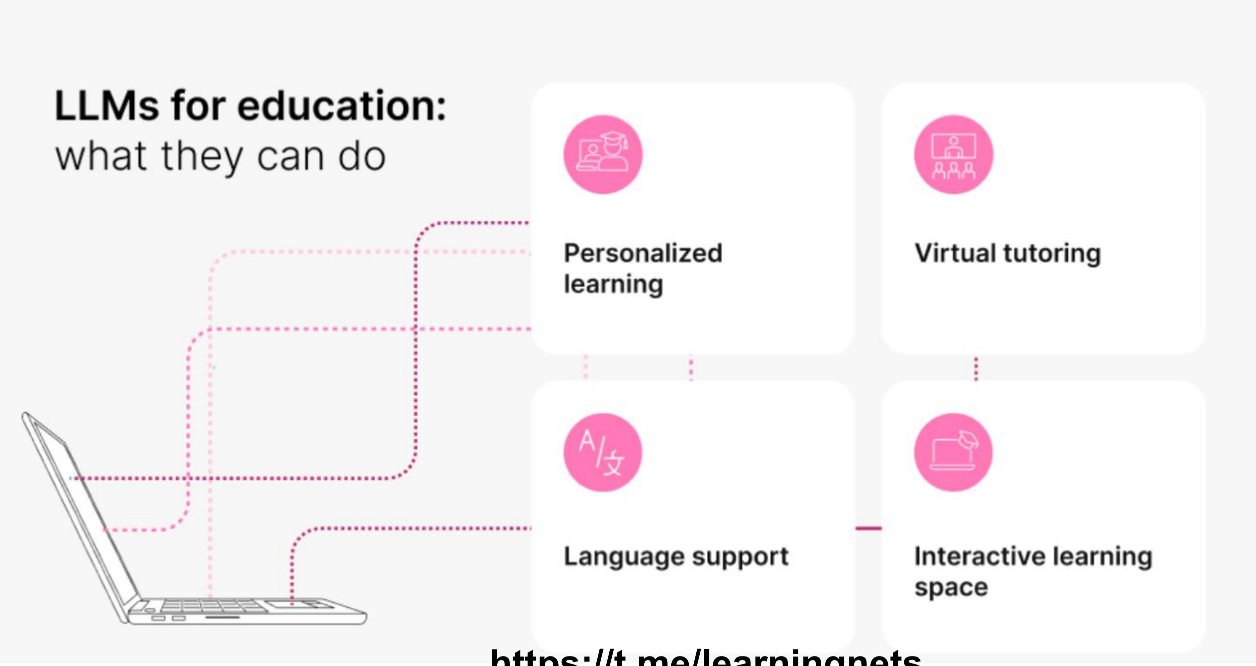

Source: pixelplex.io

AI Network Design Use Cases - Large Language Models for Network Design Insights and Automation

LLMs can provide ongoing design validation, too.

Missing route summarization in BGP or recommend better placement of redundant links in a leaf/spine architecture.

.

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Pattern recognition is a very important tool for identifying network bottlenecks in AI.

Pattern recognition involves analyzing telemetry data, logs, and flow records to identify recurring behaviors or anomalies.

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

One of the most common use cases for pattern recognition is traffic analysis.

Unbalanced traffic can lead to congestion in some areas while others remain underutilized.

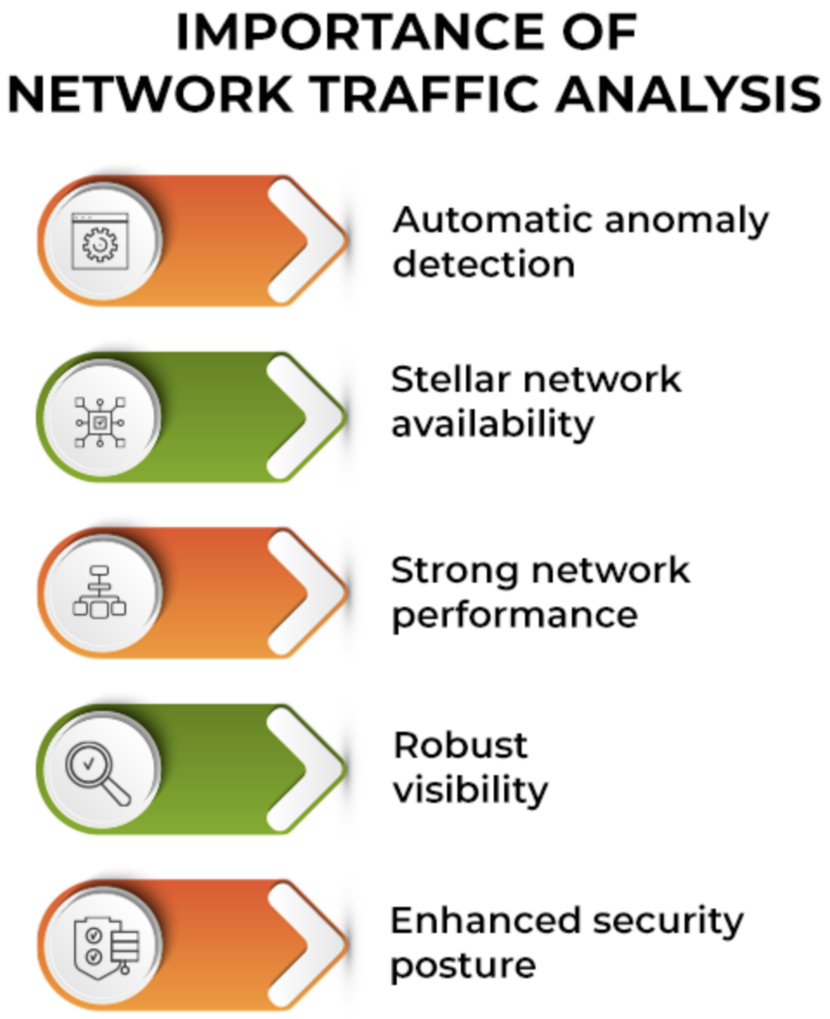

Source: spiceworks.com

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Another important area is fault detection.

Patterns such as intermittent link failures or increasing error rates often give us an idea about major outages.

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Pattern recognition systems detect early signs across devices and time, enabling proactive fixes.

For example, frequent CRC errors on a switch port may indicate a failing cable or hardware issue.

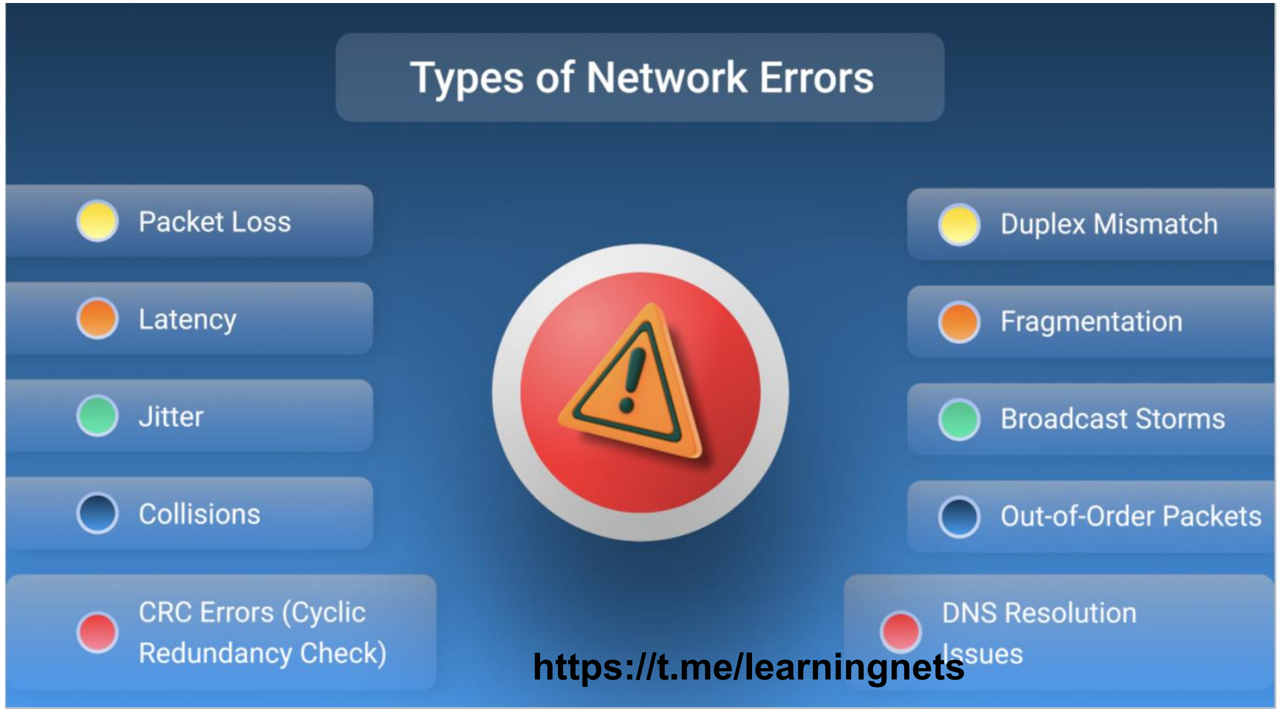

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Pattern recognition helps in capacity planning by analyzing historical data to predict when resources like CPU, memory, or bandwidth will reach their limits

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Pattern recognition detects traffic anomalies like unexpected spikes, signaling potential security breaches.

These alerts enable quicker responses and help mitigate threats effectively

AI Network Design Use Cases - Pattern Recognition for Identifying Network Bottlenecks and Optimization Opportunities

Pattern recognition improves network efficiency by identifying bottlenecks, faults, and optimization opportunities.

Its real-time data processing is critical for modern network operations and planning

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

AI and ML enhance network capabilities but require specific infrastructure based on the use case.

Their demands vary across technologies like machine learning, deep learning, large language models, and generative AI.

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Machine learning uses fewer resources than deep learning but still needs adequate compute, storage, and network capacity for data ingestion, training, and inference.

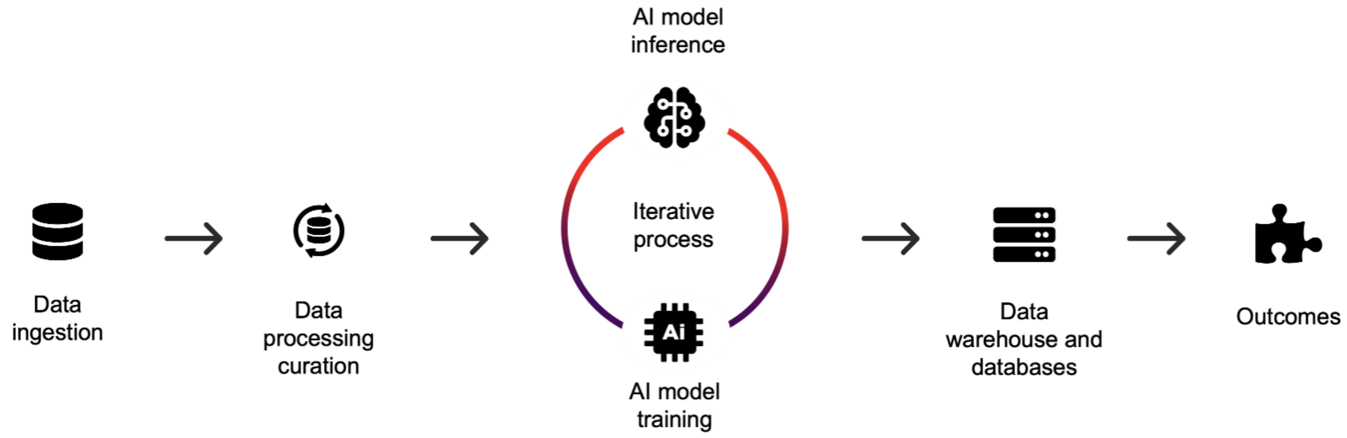

AI Workflow

Source: Equinix

https://t.me/learningnets

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Training and validation require large historical data, making fast storage systems crucial for handling structured data like logs, metrics, or NetFlow records.

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

ML workflows create heavy data traffic in data centers, especially during training compared to inference.

Sufficient bandwidth is key to ensure real-time data transfer between storage and compute nodes without delays.

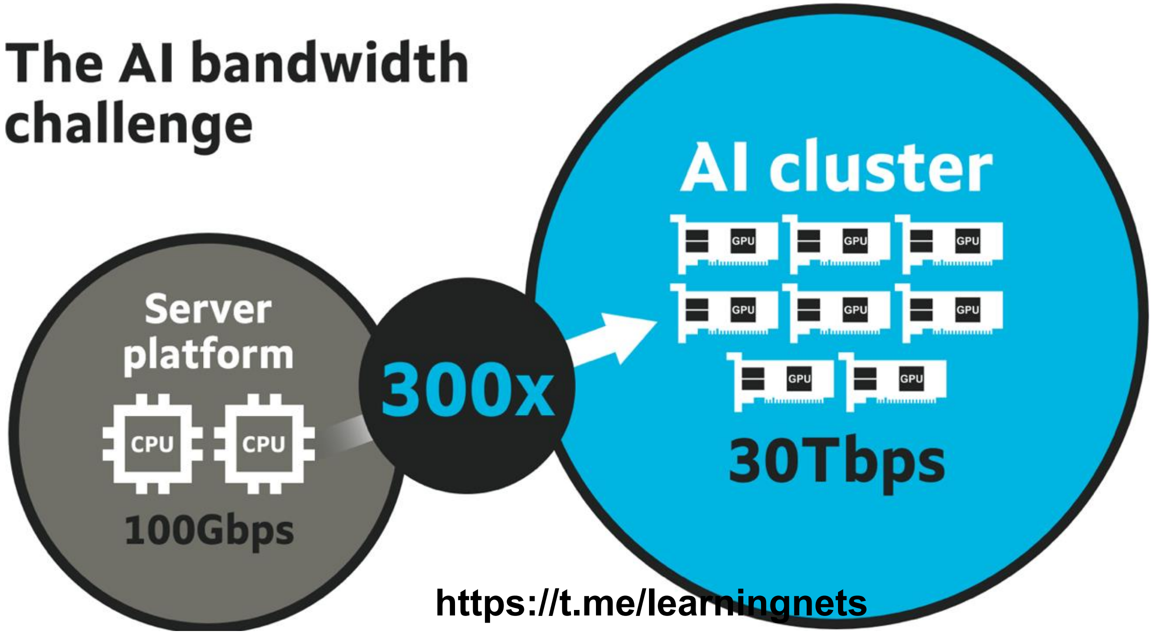

Source: Marvell

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Predictive network modeling, anomaly detection, and traffic forecasting are examples of ML-driven applications in networking.

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Deep learning, with its multi-layered neural networks, needs far more resources than regular ML.

It's ideal for tasks like image recognition or natural language processing that handle large-scale unstructured data

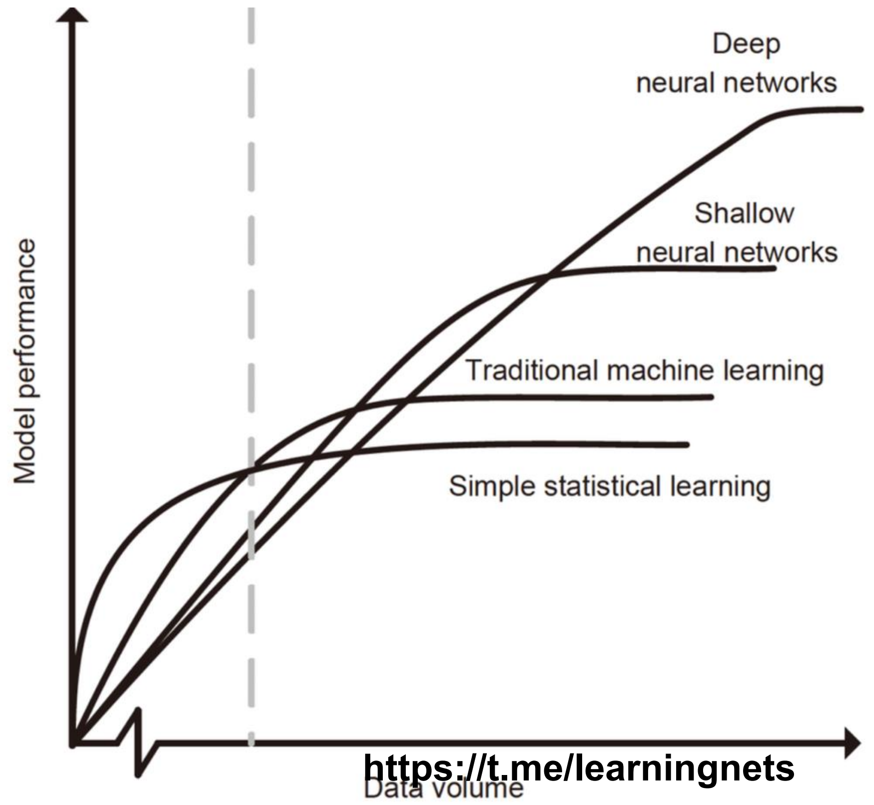

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Deep learning networks need a lot of energy.

As we covered before in power and cooling needs, data centers require advanced cooling systems to handle this!

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Large language models like GPTs require a lot of resources because of their size and training steps.

They help create insights, answer queries, and automate text tasks in networking.

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Training LLMs require high-performance clusters with GPUs or TPUs.

Inference (Using the model) also requires significant compute power, but it is less intense than training.

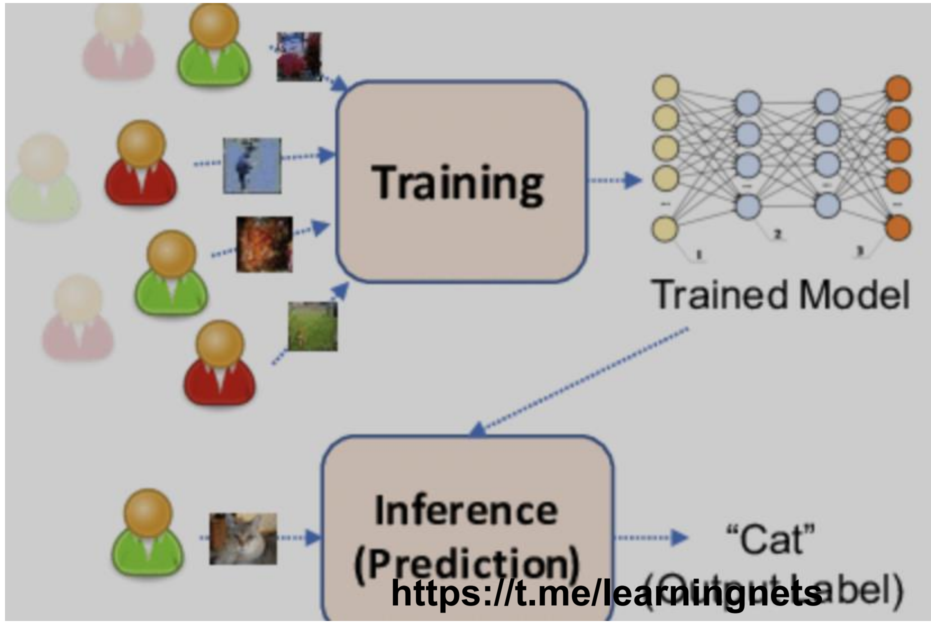

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

High-throughput, low-latency networks are needed for parameter synchronization during LLM training.

For tasks like troubleshooting or automation, LLMs require real-time access to relevant data

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

GenAI, a subset of deep learning, focuses on creating new data instead of analyzing it.

GenAI workloads use GPUs or TPUs for tasks like generating synthetic network traffic patterns

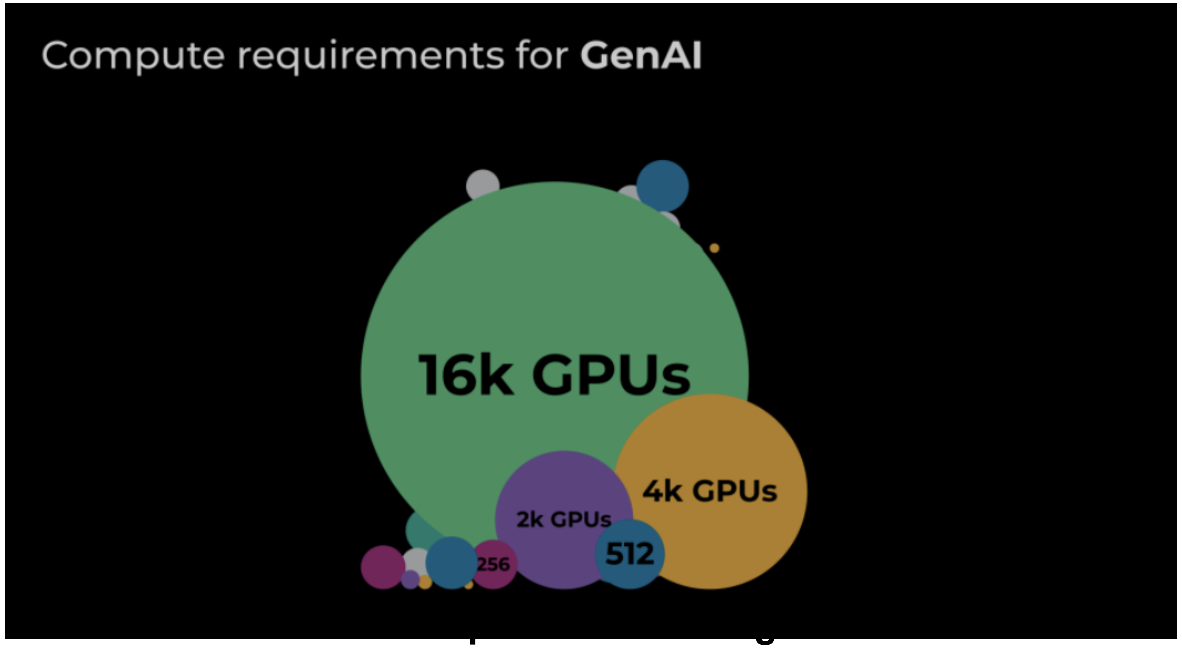

Source: Meta

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

AI workloads consume a lot of energy, needing efficient power management and cooling systems to handle the heat from GPUs and TPUs.

Source: appypie

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

High-throughput, low-latency networks are necessary for distributed AI workloads, and technologies like RDMA, InfiniBand, and NVMe-over-Fabric play a crucial role.

AI Network Design Use Cases - Impacts of infrastructure resources and requirements for different AI/ML use cases

Summary:

The infrastructure needs for AI and ML depend on the application, whether it’s machine learning, deep learning, large language models, or generative AI.

Networking for AI

https://t.me/learningnets

4. Networking for AI

High-Performance Networking Technologies

This section focuses on technologies designed to enable low-latency, high-bandwidth communication for AI workloads.

RDMA
InfiniBand
RoCE
RoCEv2
iWARP
Ultra Ethernet
UEC

Compute Accelerators and Interconnects

This section delves into the hardware and interconnects that power AI computation and facilitate seamless communication between compute nodes.

DPU, TPU, GPU
NVLink
NVMe-over-Fabrics (NVMe-oF)

4. Networking for AI

Traffic Management and Congestion Control

This section focuses on techniques and technologies used to manage and optimize network traffic for AI workloads, ensuring reliable and efficient communication.

Layer 2

DCB – Datacenter Bridging

- **PFC (Priority Flow Control)**
- **ETS (Enhanced Transmission Selection)**
- **QCN (Quantized Congestion Notification)**
- **DCBX (Datacenter Bridging Exchange)**

Layer 3

ECN (Explicit Congestion Notification)

DCQCN (Data Center Quantized Congestion Notification)

TIMELY

HPCC (High Precision Congestion Control)

RDMA lets one computer directly access another's memory without using the operating system or CPU.

It speeds up communication by bypassing the CPU.

Normally, the CPU handles data transfers over a network, which takes extra time.

With RDMA, no CPU is involved, data goes directly to the memory of the other computer.

This reduces latency so the CPU can handle other important tasks.

.

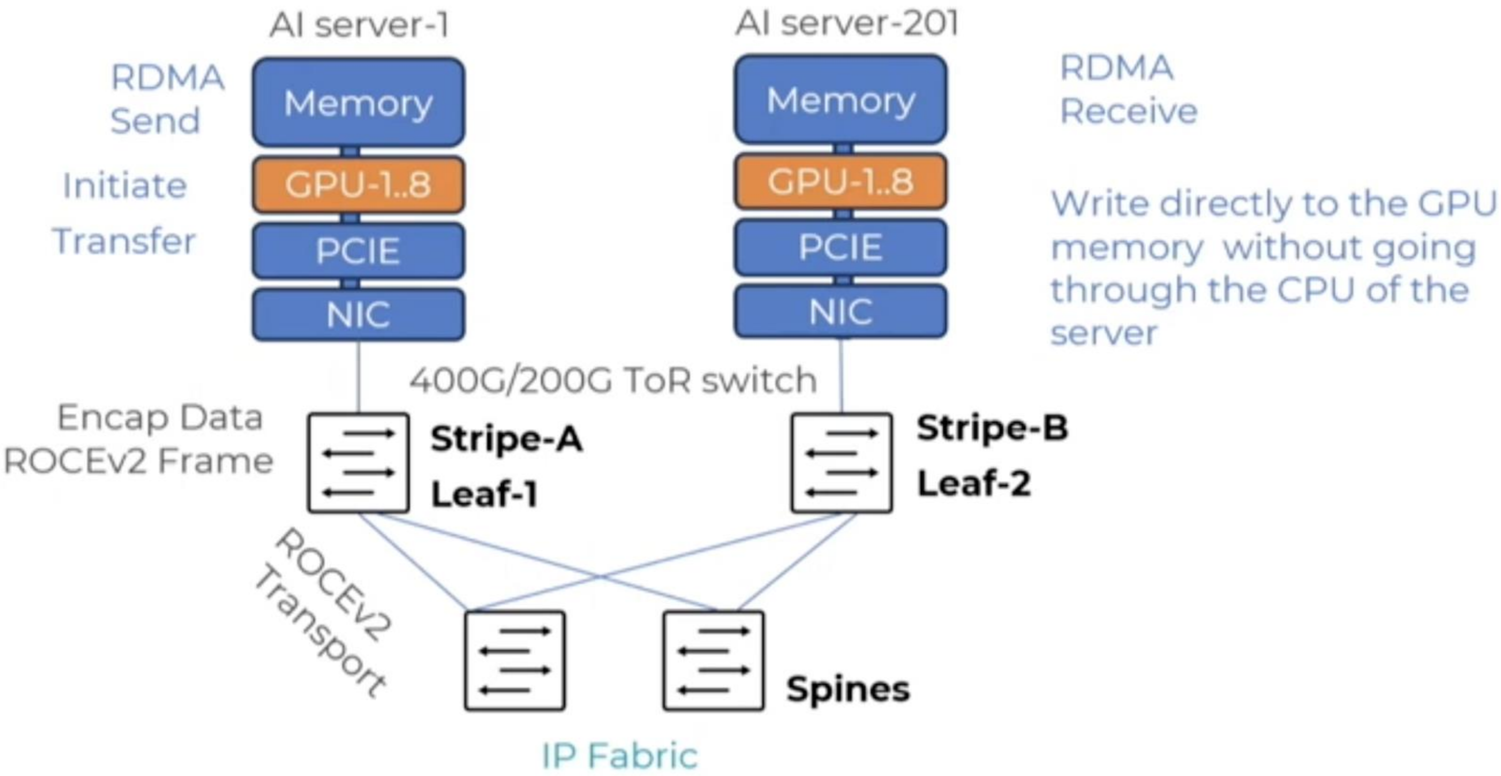

RDMA is especially useful in environments where speed is critical, like data centers and AI workloads.

For example, in AI training, large amounts of data must be quickly shared between GPUs or TPUs.

.

Another benefit of RDMA is that it reduces the amount of work the CPU has to do.

CPU bypass, RDMA saves energy and improves the system's overall performance.

.

In AI, RDMA is ideal because AI workloads need parallel computing across nodes like GPUs, TPUs, or DPUs.

.

RDMA performs best with high-performance networks like InfiniBand or RoCE, designed to support its needs (we'll cover these later).

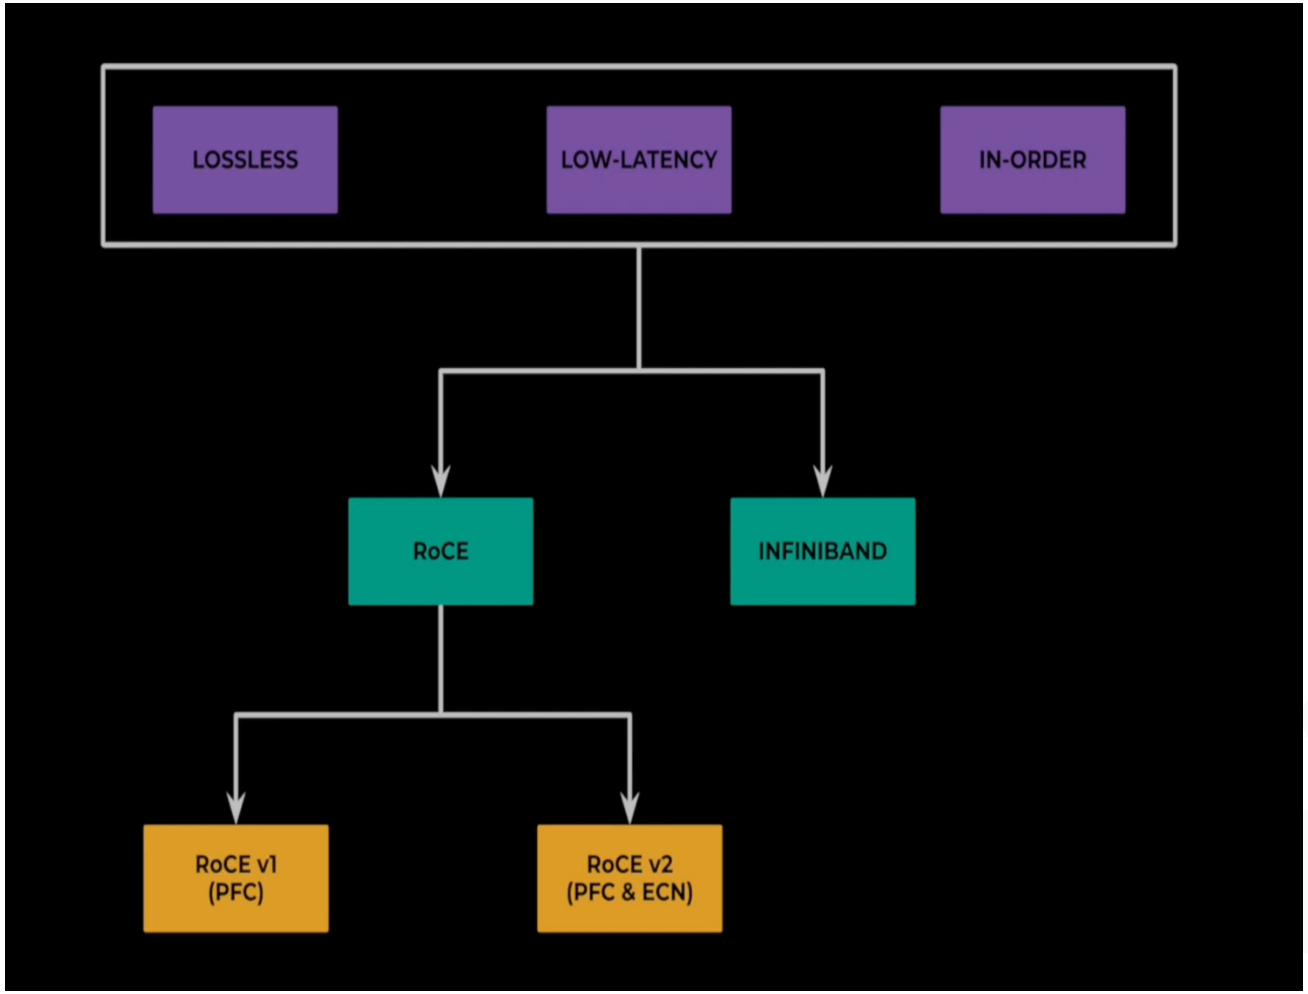

These protocols deliver the low latency and high bandwidth RDMA requires to work efficiently.

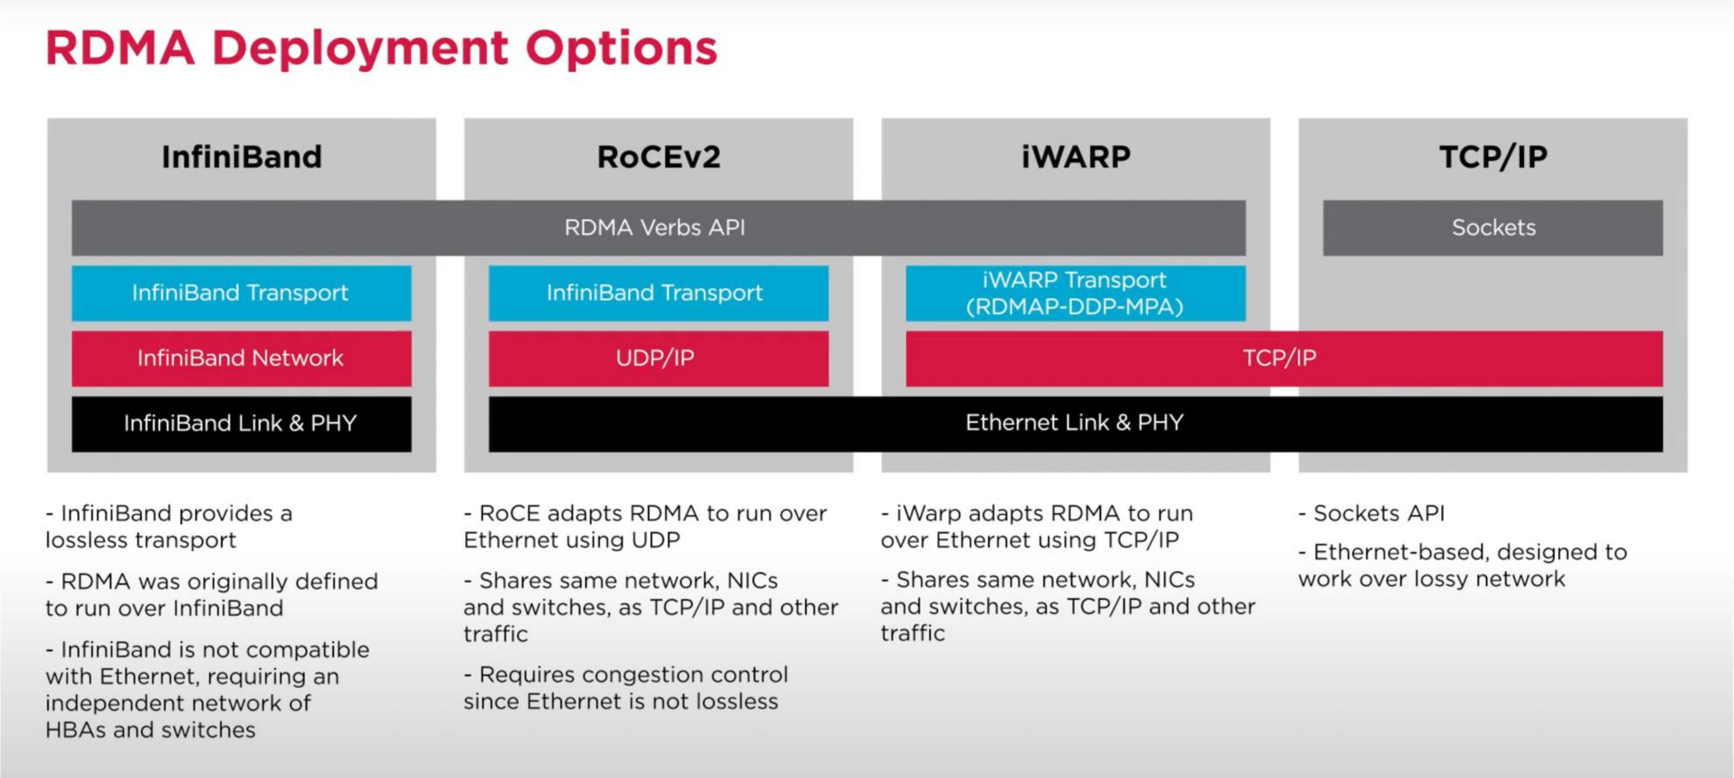

Source: Broadcom

.

Like any protocol, RDMA has its challenges.

It needs specialized hardware like RNICs and high-performance networks like InfiniBand or RoCE, which add costs.

Congestion management is another issue, handled by methods like ECN and DCQCN, which we’ll discuss later.

Overhead with TCP/IP due to the multiple layers of networking stack.

TCP retransmissions, and CPU involvement in packet handling.

RDMA eliminates overhead by bypassing the kernel and CPU and directly accessing memory.

Summary:

RDMA is a fundamental technology for AI.

Deliver low-latency, high-bandwidth communication, thus best for AI workloads.

.

Infiniband

Infiniband is designed specifically for environments where speed and low latency are critical.

Unlike traditional Ethernet networks, InfiniBand offers ultra-low latency and extremely high bandwidth, ideal for HPC, supercomputers, and AI training clusters.

Infiniband

It’s built to deliver direct memory-to-memory data transfers between nodes, bypassing the operating system and CPU entirely.

InfiniBand networks are often used in distributed AI training

The high throughput and low latency allow AI models to train faster, which is crucial as models become larger and require more computing power.

Infiniband

InfiniBand uses a credit-based flow control system to manage congestion

Fiber channel, used in storage area networks, also work based on credit-based flow control.

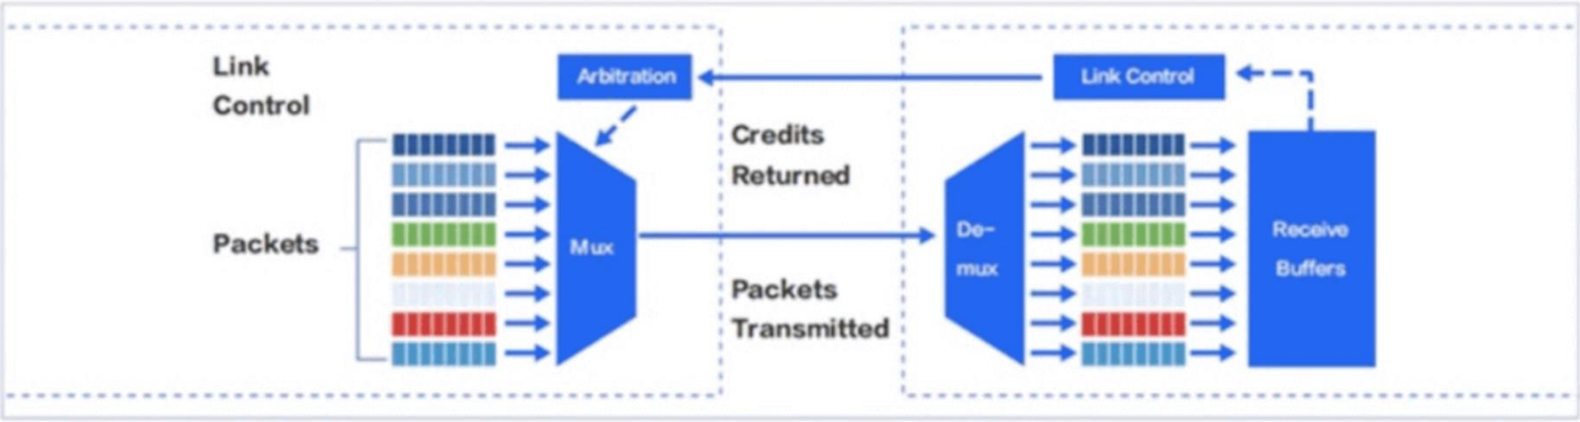

Infiniband

From a hardware perspective, InfiniBand requires specialized adapters (HCA), switches, and cables.

More expensive and less flexible compared to Ethernet-based solutions.

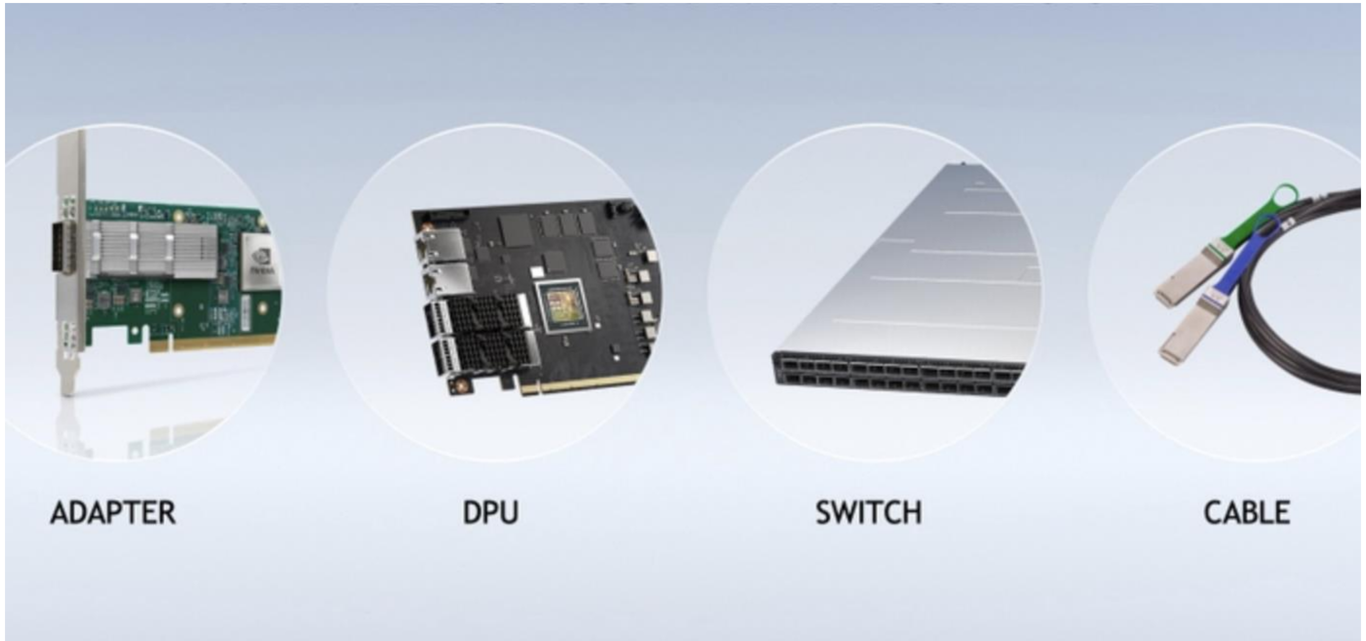

Source: Nvidia

Infiniband

InfiniBand is tightly integrated with RDMA, meaning it supports direct memory access over its network.

Infiniband

Infiniband operates at Layer 2, meaning it doesn’t use IP addressing.

Instead, it works with its own addressing system.

Infiniband

One of the main advantages of InfiniBand is its low latency

.

Infiniband

InfiniBand currently supporting up to 200 Gbps with HDR and 400 Gbps with NDR as you can see from the table below.

Generation	Transport Speed per Lane	Total Speed (x4 Link)
SDR (Single Data Rate)	2.5 Gbps	10 Gbps
DDR (Double Data Rate)	5 Gbps	20 Gbps
QDR (Quad Data Rate)	10 Gbps	40 Gbps
FDR (Fourteen Data Rate)	14 Gbps	56 Gbps
EDR (Enhanced Data Rate)	25 Gbps	100 Gbps
HDR (High Data Rate)	50 Gbps	200 Gbps
NDR (Next Data Rate)	100 Gbps	400 Gbps

.

Infiniband

InfiniBand currently supporting up to 200 Gbps with HDR and 400 Gbps with NDR as you can see from the table below.

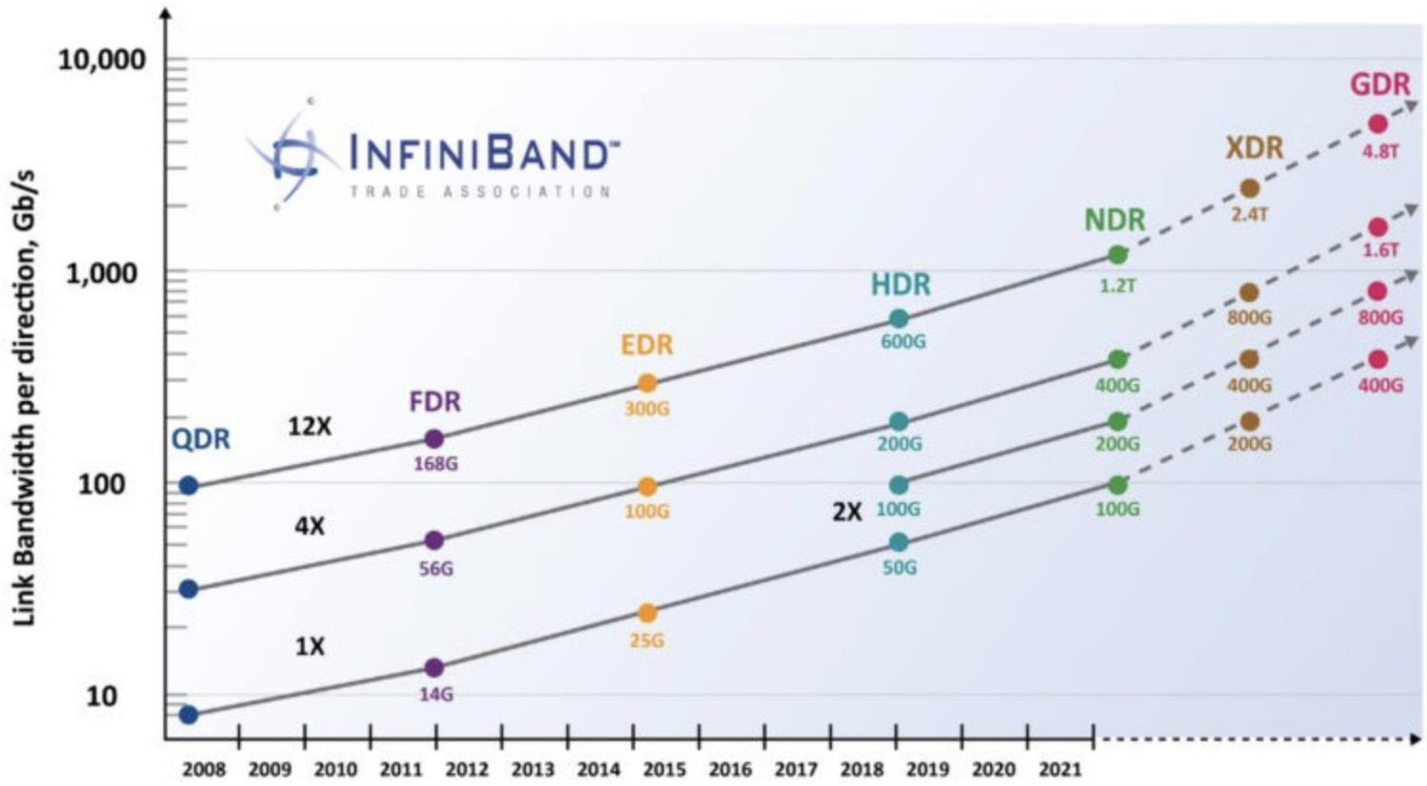

.

https://t.me/learningnets

Source: IBTA

Infiniband

InfiniBand has some drawbacks.

Requires specialized hardware like HCAs, which are Host channel adapters, similar to the Ethernet NICs, but it's preparatory.

Also, preparatory Infiniband switches are more expensive than Ethernet equipment.

Infiniband

Summary:

InfiniBand is a Layer 2, high-performance technology designed for speed, scalability, and efficiency.

Unlike open-standard Ethernet, it is proprietary, making it ideal for AI workloads and HPC environments where low latency and high throughput are critical.

However, its cost and complexity limit its use in general networking.

High Performance Networking Technologies - RoCE

RoCE, or RDMA over Converged Ethernet, is a protocol that allows RDMA to work over Ethernet.

RoCE brings the benefits of RDMA to Ethernet, which is a widely used networking technology.

RoCE operates over Layer 2 of the OSI model

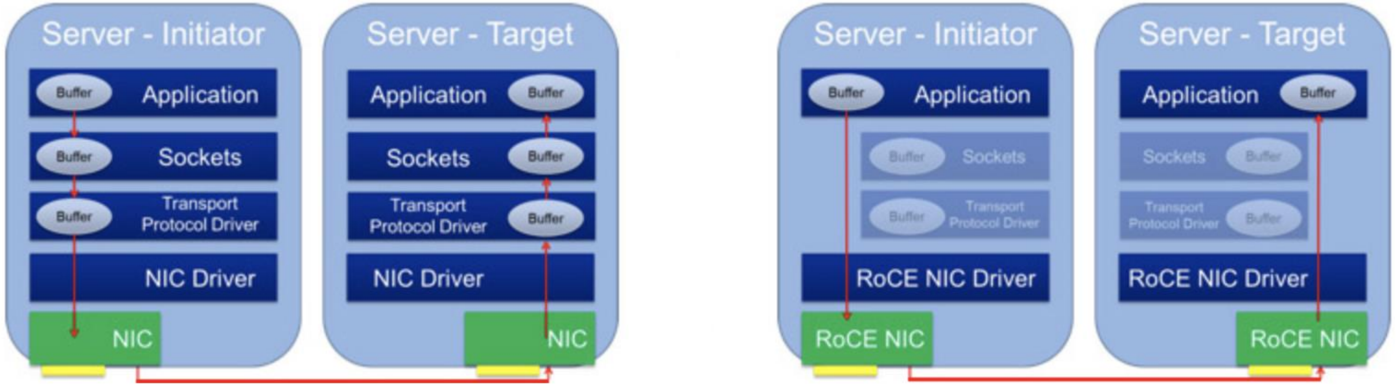

Traditional vs RoCE Data communication

High Performance Networking Technologies - RoCE

One of the major advantages of RoCE is its low latency, making it ideal for distributed applications such as AI training or storage transfers.

Message Size	Average Latency		
	Native IB (us)	*RoCE (us)*	*Native Ethernet (us)*
8	0.892	0.988	12.662
16	0.938	1.014	12.692
32	0.95	1.026	12.718
64	0.958	1.024	12.712
128	1.344	1.418	12.78
256	1.402	1.474	12.884
512	1.518	1.584	13.104
1024	1.752	1.814	13.55

.

Source: Naddod

High Performance Networking Technologies - RoCE

RoCE is also more cost-effective compared to specialized solutions like InfiniBand, as it uses standard Ethernet hardware.

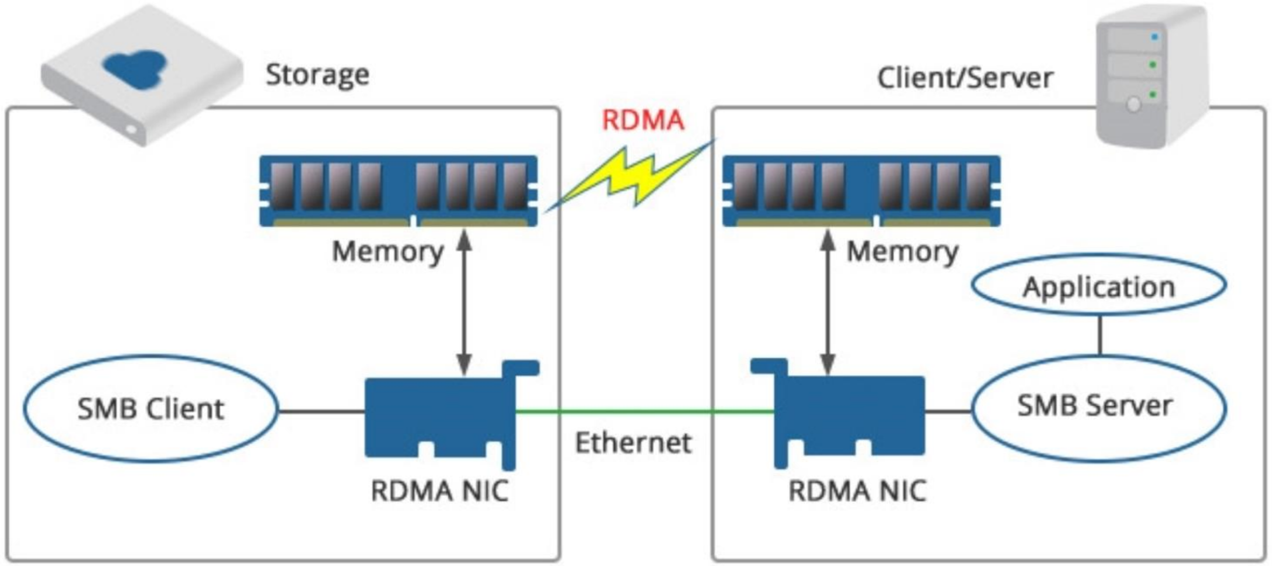

High Performance Networking Technologies - RoCE

RoCE has its challenges too.

Because it relies on Ethernet, it introduces potential issues with congestion control.

Without proper congestion control, packet loss can happen

Techniques like ECN, Explicit Congestion Notification and PFC are often necessary to address this.Specifically PFC and other DCB suites.

High Performance Networking Technologies – RoCEv2

RoCEv2 is an enhanced version of RoCE, brings Layer 3 support.

While RoCE operates only at Layer 2, RoCEv2 works across Layer 3 networks

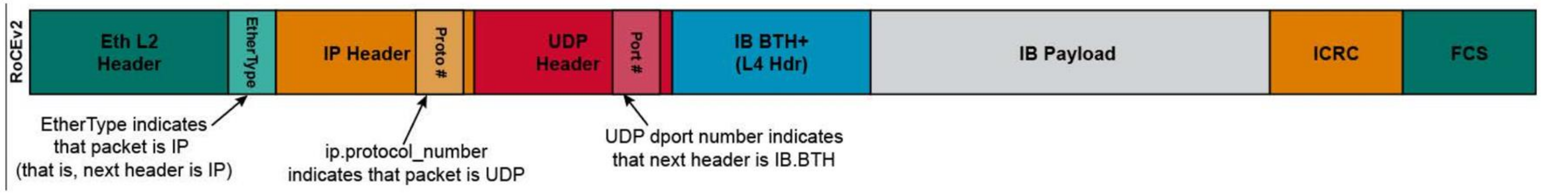

RoCEv2 Header

https://t.me/learningnets

.

High Performance Networking Technologies – RoCEv2

RoCEv2 supports IP addressing, enabling RDMA communication over Layer 3 TCP/IP networks.

This allows RDMA to work in complex topologies, including inter-data center communication, unlike RoCE

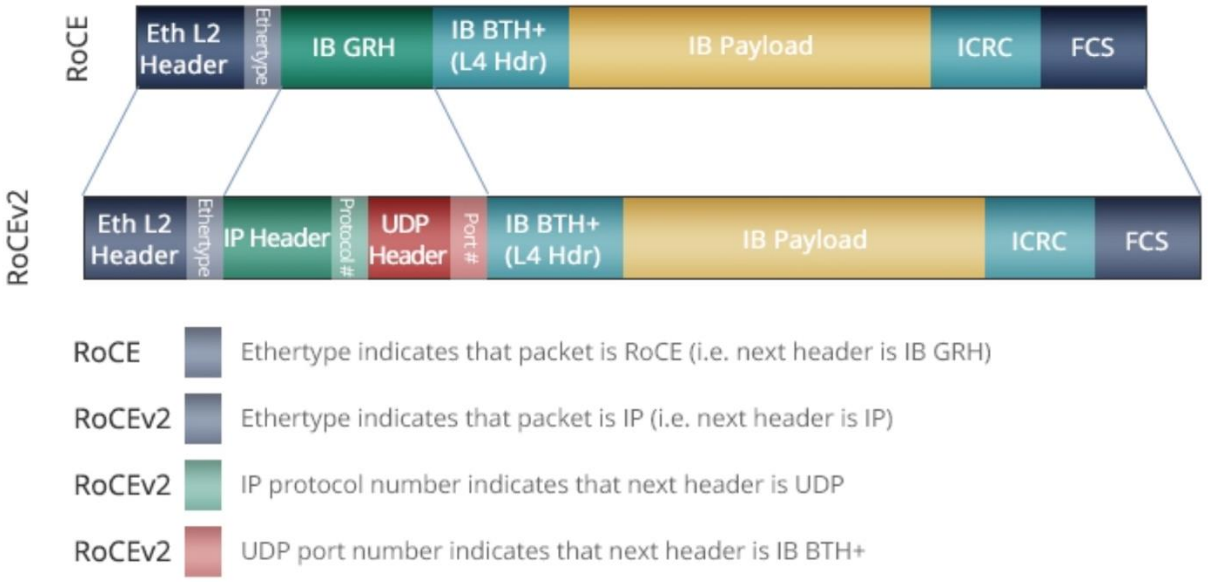

.

RoCE and RoCEv2 Headers

High Performance Networking Technologies – RoCEv2

RoCEv2 can integrate seamlessly with standard IP networking, making it more flexible than InfiniBand or RoCE.

Remember, Infiniband is Layer 2 as well.

High Performance Networking Technologies – RoCEv2

RoCEv2 has slightly more overhead than RoCE due to operating over Layer 3 with IP headers.

This can affect ultra-low-latency performance, but the impact is usually minor.

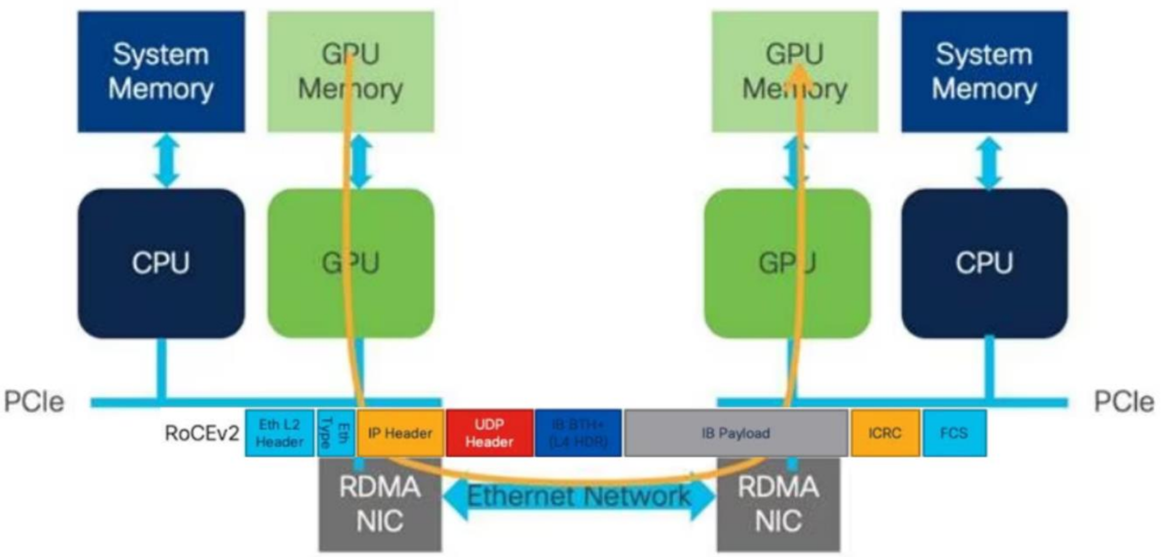

RoCEv2 Communication

https://t.me/learningnets

High Performance Networking Technologies – RoCEv2

For the best performance, it relies on ECN-enabled switches, and without ECN, performance can degrade due to packet loss.

High Performance Networking Technologies – RoCEv2

Classical Ethernet switches might struggle with RoCE or RoCEv2 since these protocols need advanced features like ECN, DCQCN, PFC, Timely, and HPCC.

We'll discuss these in detail later.

High Performance Networking Technologies – RoCEv2

RoCE and RoCEv2 depend on lossless communication for RDMA.

Efforts exist to support RDMA over lossy networks, but packet loss is problematic since RDMA lacks retransmission mechanisms.

High Performance Networking Technologies – RoCEv2

RoCE and RoCEv2 require Priority Flow Control (PFC) to avoid packet loss, a feature found in DCB-enabled Ethernet switches.

Classical Ethernet switches lack PFC, making them unsuitable for RDMA traffic.

High Performance Networking Technologies – RoCEv2

RoCEv2, in particular, mitigates this issue by using Explicit Congestion Notification, ECN and Data Center Quantized Congestion Notification, DCQCN, which work together to detect and manage congestion dynamically.

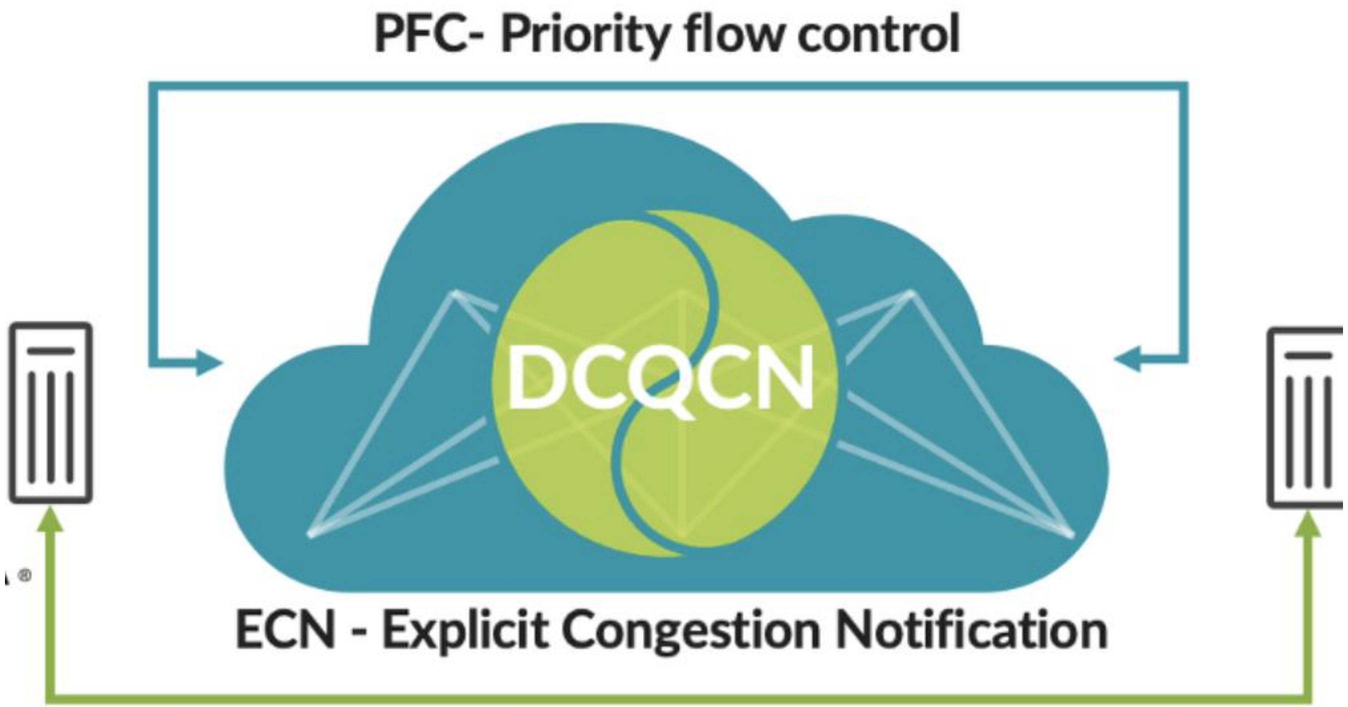

.

Source: NVdia

https://t.me/learningnets

High Performance Networking Technologies – RoCEv2

Summary:

Classical Ethernet switches lack the lossless, low-latency, and high-performance capabilities required for RoCE or RoCEv2.

Features like PFC, ECN, QoS, and Layer 3 are needed for effective use.

High Performance Networking Technologies – iWARP

iWARP is a networking standard that brings RDMA functionality to traditional TCP/IP-based networks.

.

agenda

High Performance Networking Technologies – iWARP

Unlike InfiniBand or RoCE, which require specialized hardware, iWARP works over standard Ethernet and IP infrastructures.

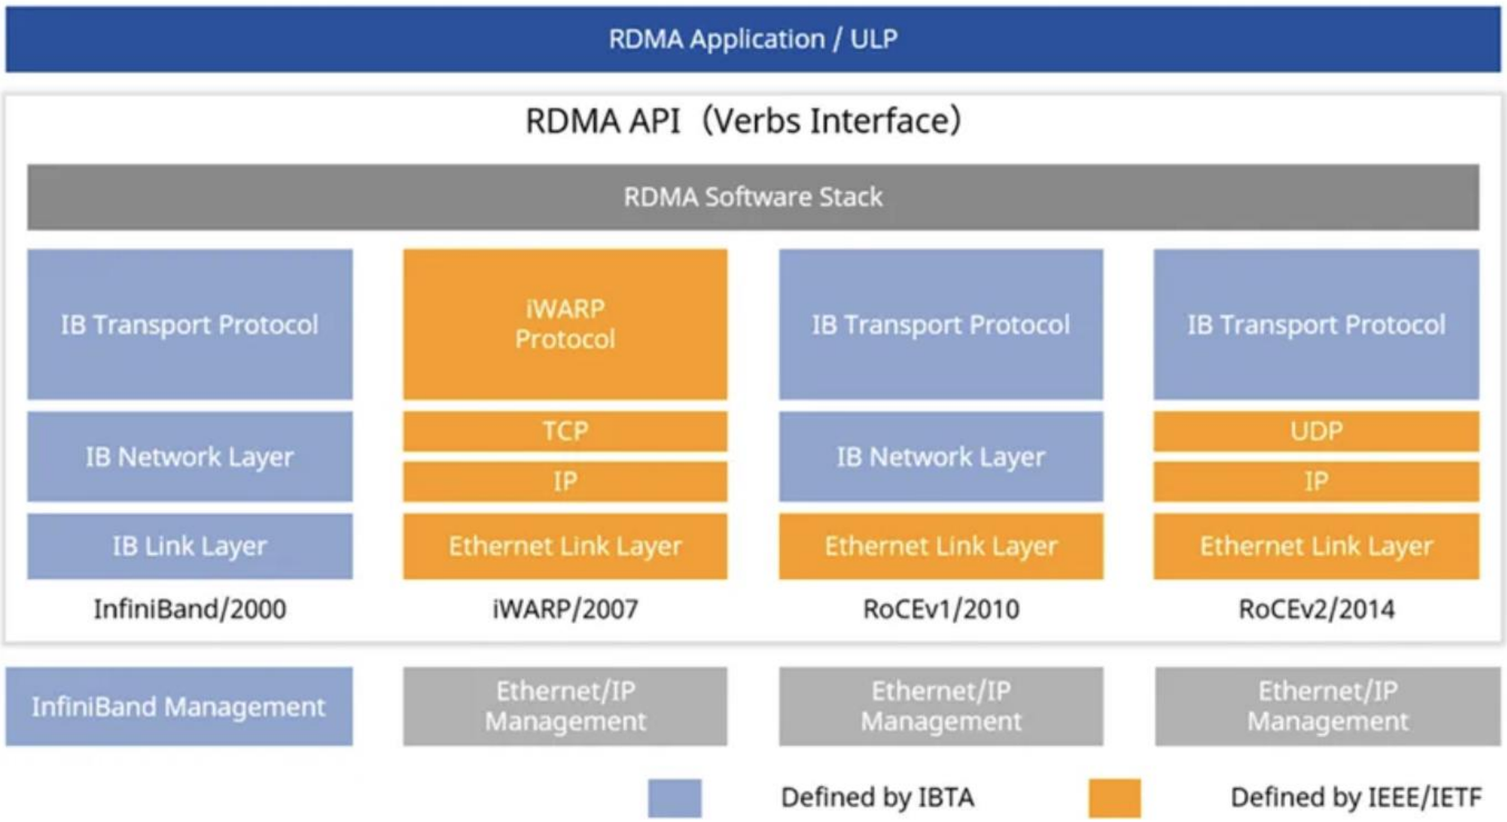

Source: Medium

High Performance Networking Technologies – iWARP

It runs RDMA over TCP/IP, which makes it widely deployable and easier to integrate.

iWARP is beneficial in environments where hardware or network upgrades for InfiniBand or RoCE are not feasible.

High Performance Networking Technologies – iWARP

If a company wants RDMA but only has classical Ethernet switches without PFC or ECN, iWARP enables RDMA without requiring upgrades.

.

High Performance Networking Technologies – iWARP

Trade-offs using iWARP:

Using TCP/IP introduces latency and overhead compared to RoCE or InfiniBand.

While TCP improves reliability through retransmissions, this added latency can slow down real-time AI or high-frequency trading.

High Performance Networking Technologies – iWARP

For applications that reliability and ease of deployment are more important than latency, iWARP is a strong choice

.

High Performance Networking Technologies – iWARP

iWARP offers decent RDMA performance but lacks the ultra-low latency of InfiniBand or RoCE, making it less ideal for tasks like LLM training.

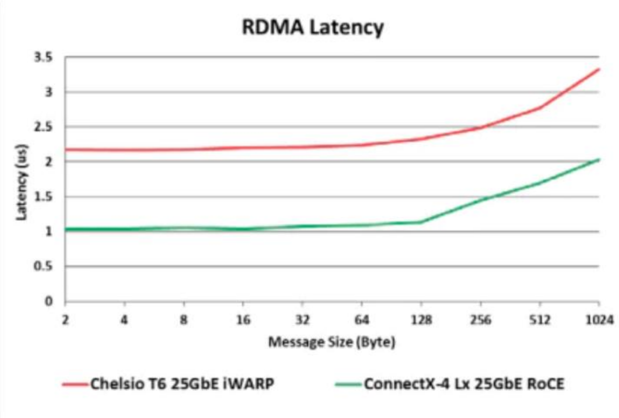

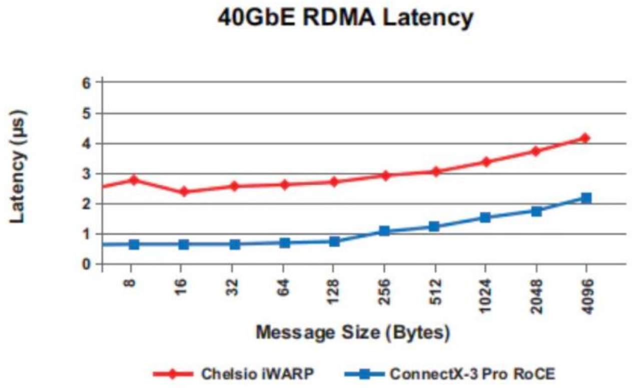

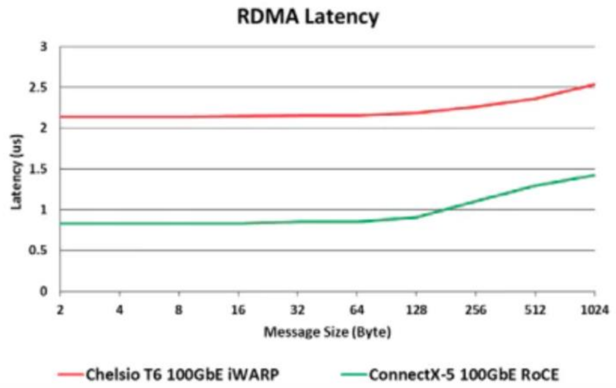

25G, 40G, 100G Latency Test by Naddod
RoCE is much faster than iWARP
Source: Naddod

.

High Performance Networking Technologies – iWARP

iWARP has loss tolerance.

So, it can tolerate packet loss because it runs over TCP, but RoCE and RoCEv2 rely on lossless Ethernet.

.

High Performance Networking Technologies – iWARP

Summary:

Although it doesn't deliver the low latency as InfiniBand or RoCE, it provides reliability and interoperability.

Ethernet

Ultra Ethernet is a new networking technology designed for AI and HPC needs.

It handles massive traffic, low latency, and high bandwidth, unlike classical Ethernet, which is general-purpose.

Ethernet

Ultra Ethernet targets to combine the broad adoption of Ethernet with the performance features of InfiniBand.

Ultra Ethernet offers advanced congestion control, flow mechanisms, and scalability, ideal for AI environments.

Ethernet

Ultra Ethernet might also face challenges.

Classical Ethernet's decades-old ecosystem is vast, with devices, software, and expertise.

Transitioning will need new hardware like switches, NICs, and software updates.

Ethernet

Interoperability could be an issue.

While Ethernet is widely compatible, Ultra Ethernet’s new features might not work smoothly with older devices.

However, it is expected to remain mostly compatible with Classical Ethernet.

Ethernet

Ultra Ethernet is cheaper than InfiniBand but still costs more than traditional Ethernet hardware.

Ultra Ethernet is still new, so its ecosystem isn’t as mature as classical Ethernet or InfiniBand yet.

Advanced congestion management and flow control require specialized hardware and expertise, so the solution might be more complex compared to classical Ethernet.

Ultra Ethernet doesn’t yet have the same level of maturity as classical Ethernet or InfiniBand.

Ethernet

<u>Ultra ethernet vs. Classical Ethernet Comparison</u>

Classical Ethernet wasn't built for AI or HPC needs like high throughput and low latency.

It uses a best-effort approach, where congestion and packet loss are normal.

.

Ethernet

Ultra ethernet vs. Classical Ethernet Comparison

Ultra Ethernet offers congestion control, better flow control, and robust transport without altering Layer 2 semantics.

.

Ethernet

Ultra ethernet vs. Classical Ethernet Comparison

Ultra Ethernet supports lossless communication, unlike classical Ethernet.

While classical Ethernet is general-purpose and cost-effective, Ultra Ethernet is built for high-performance environments.

Ethernet

Ultra ethernet vs Infiniband and RoCE/RoCEv2

Ultra Ethernet matches InfiniBand in latency and throughput but offers broader adoption and compatibility.

InfiniBand, being proprietary, requires specific hardware and software, making it more expensive and complex.

Ethernet

Ultra ethernet vs Infiniband and RoCE/RoCEv2

As of the beginning of 2025, Ultra Ethernet is still under development.

Ethernet

<u>Ultra ethernet vs Infiniband and RoCE/RoCEv2</u>

Ultra Ethernet integrates congestion control and scalability directly into Ethernet, avoiding external extensions like RDMA over Converged Ethernet.

Ethernet Consortium

The UEC is a new standards body advancing Ultra Ethernet with best practices and hardware/software requirements.

The consortium includes industry leaders who want to enhance Ethernet to support AI requirements.

Ethernet Consortium

The UEC was founded by Intel, Broadcom, Arista, and Meta, leveraging their networking and data center expertise.

Founding Members

AMD · ARISTA · BROADCOM · CISCO · EVIDEN an atos business · Hewlett Packard Enterprise · intel · Meta · Microsoft

.

Ethernet Consortium

UEC members collaborate to design Ultra Ethernet, aiming to create high-performance Ethernet standards.

.

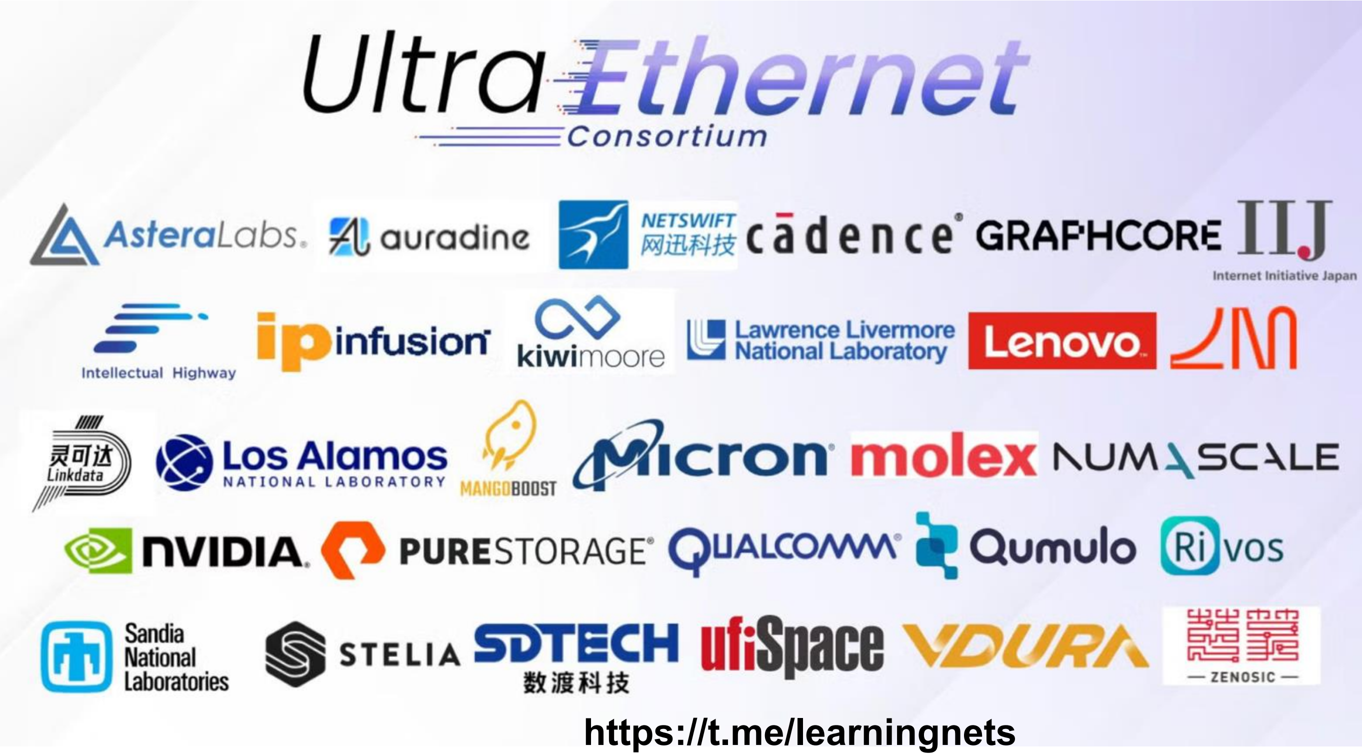

Ethernet Consortium

Currently, in 2025, the UEC is defining the technical specifications for Ultra Ethernet.

and TPU

A DPU handles data tasks, offloading work from the CPU, similar to BFD offloading processing.

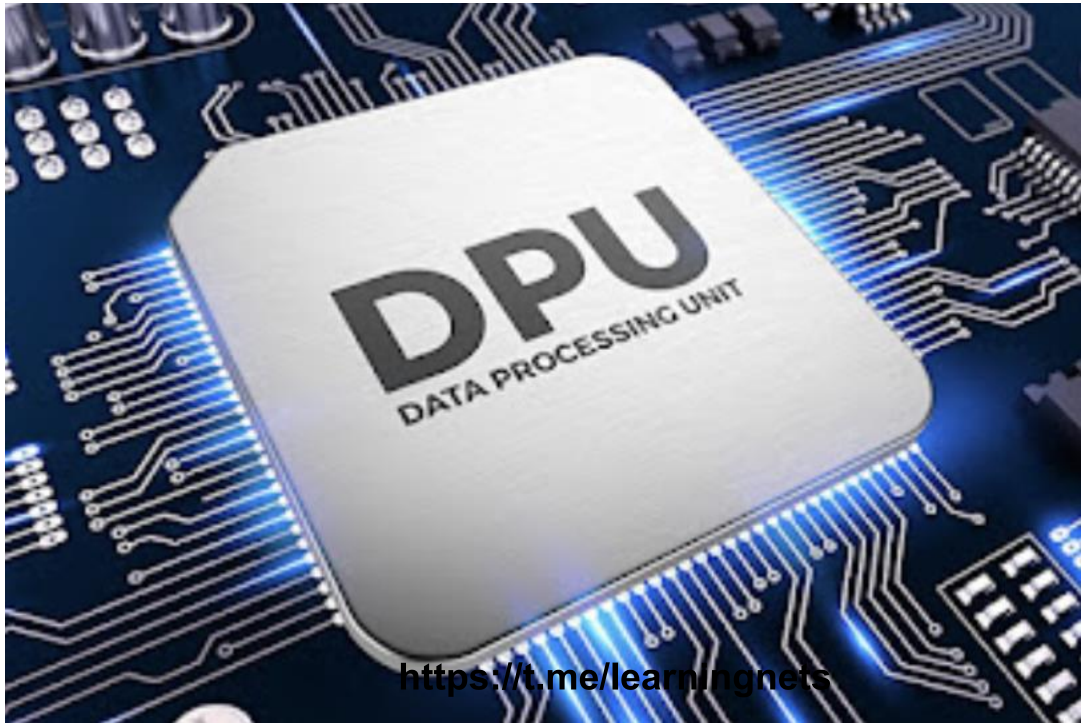

and TPU

A DPU is a smart NIC with computing power for tasks like packet processing, encryption, and storage.

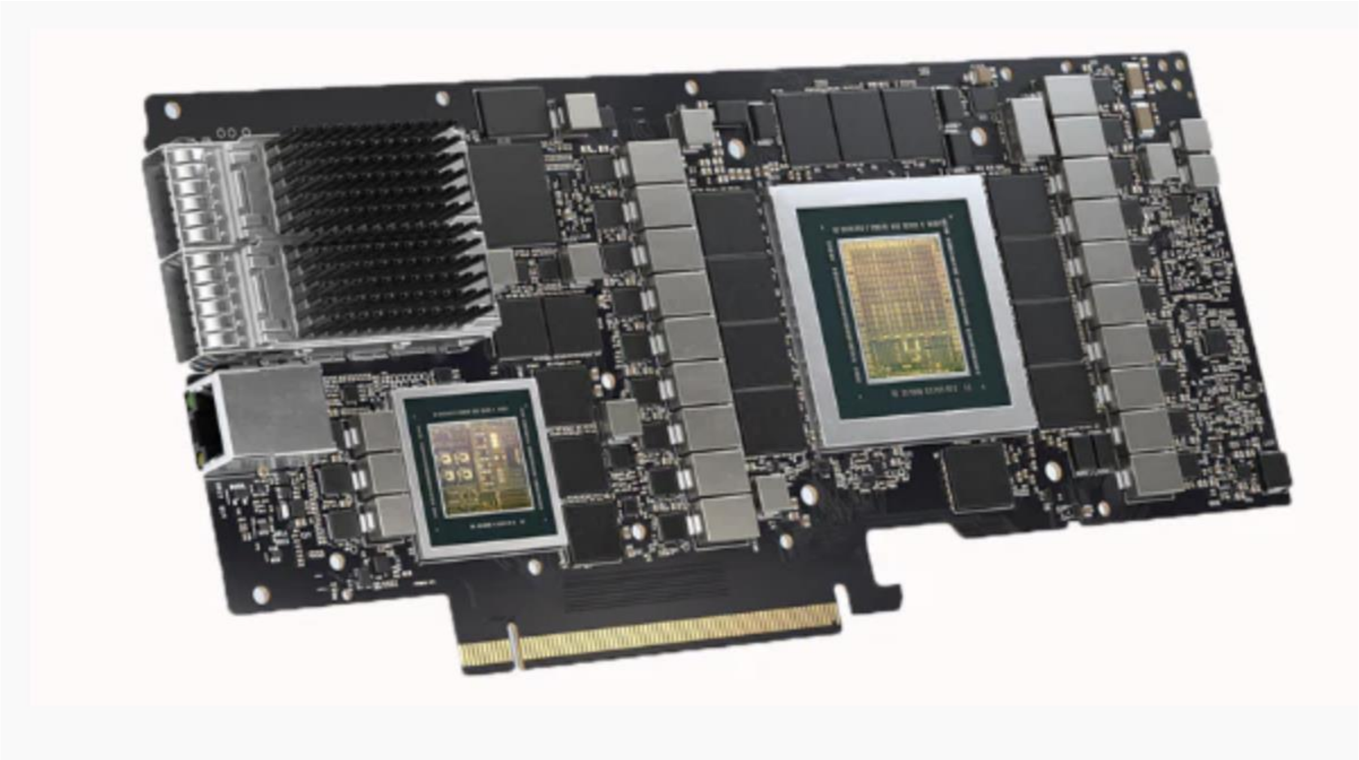

DPU - Data Processing Unit

and TPU

DPUs are essential in large-scale datacenters but they are not as good as GPU or TPU for AI tasks.

and TPU

GPUs were made for graphics but now power parallel tasks like AI and ML.

They split tasks into smaller ones, running them on thousands of cores.

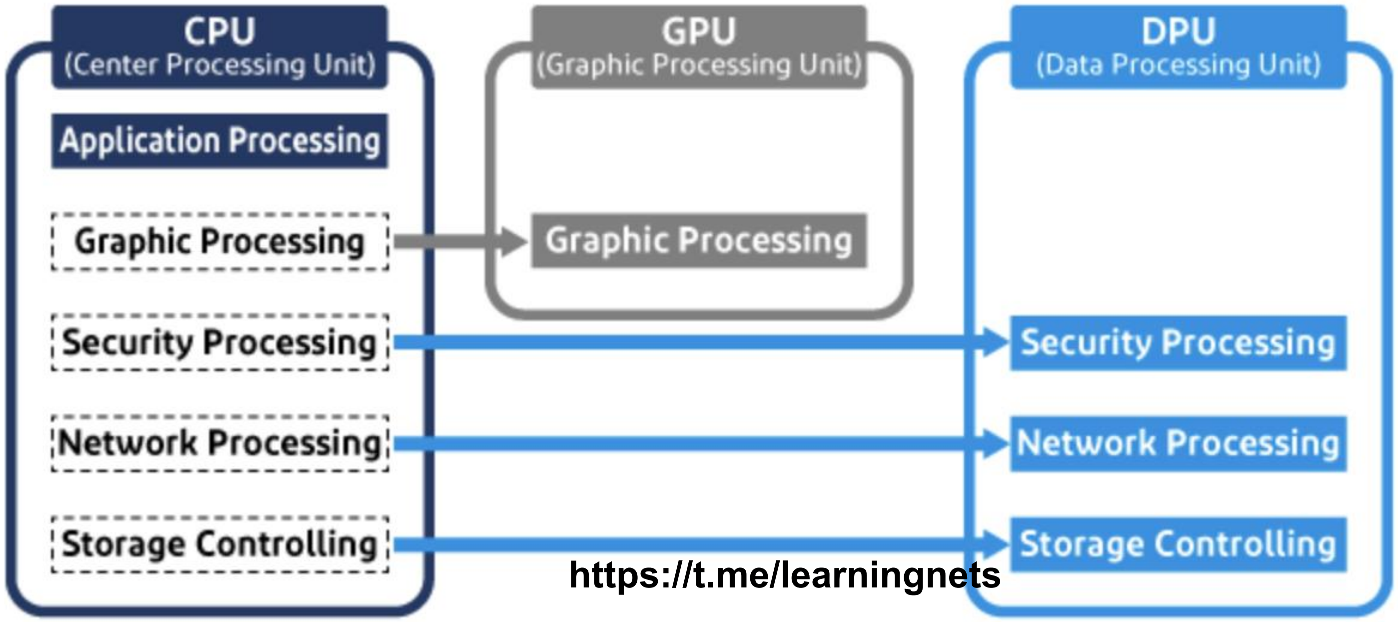

and TPU

GPUs are the primary hardware computing nodes in AI environments, particularly for training deep learning models.

and TPU

AI training processes large data in parallel, like matrix multiplications.

GPUs handle this much faster than CPUs.

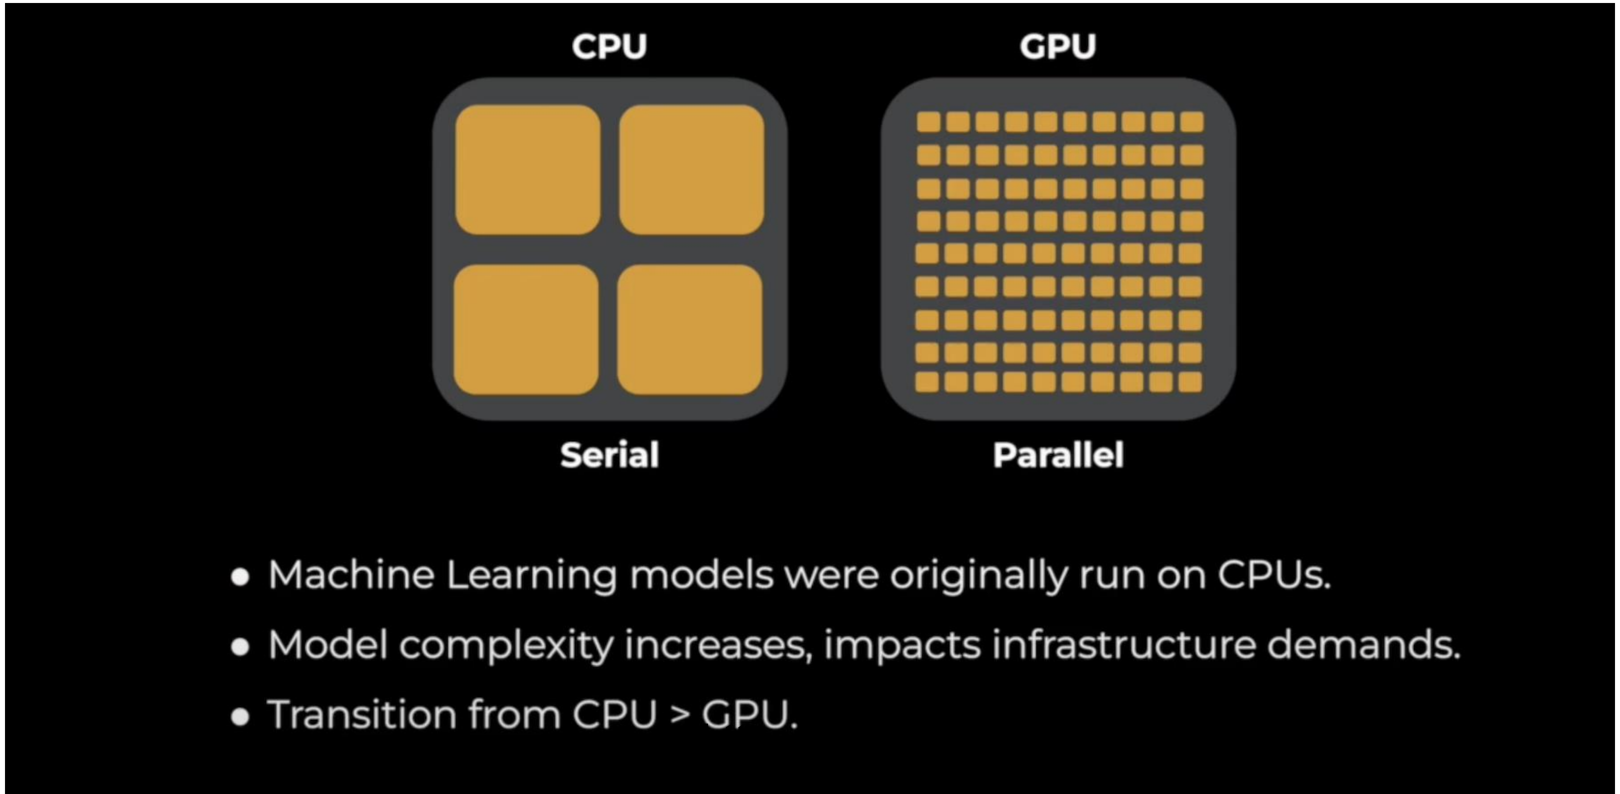

and TPU

NVIDIA with CUDA GPUs and AMD with Radeon GPUs dominate the market, with NVIDIA leading in AI due to its CUDA platform.

and TPU

TPUs are Google-designed processors for machine learning.

Unlike GPUs, TPUs focus on tensor tasks, excelling in deep learning.

and TPU

TPUs are heavily used in Google's AI services, such as Google Translate, Google Photos, and Google Search, as well as in Google Cloud Platform, GCP.

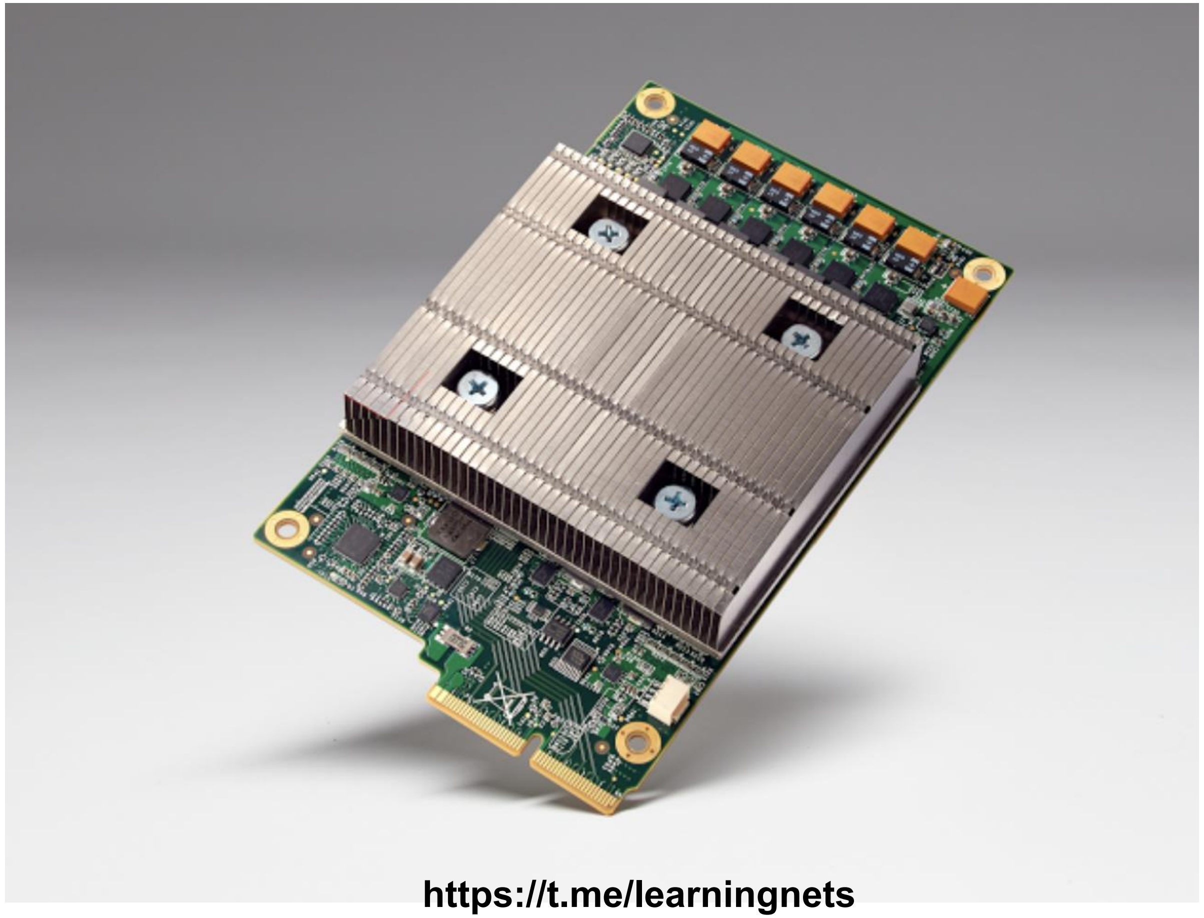

Google's TPU

and TPU

TPUs are Google’s proprietary hardware, not sold off-the-shelf.

They’re available via GCP for AI workloads, ideal in Google’s cloud.

and TPU

DPU vs GPU vs TPU:

Purpose:

DPUs handle data tasks like networking and encryption.

GPUs excel at parallel computing, while TPUs focus on AI tensor operations.

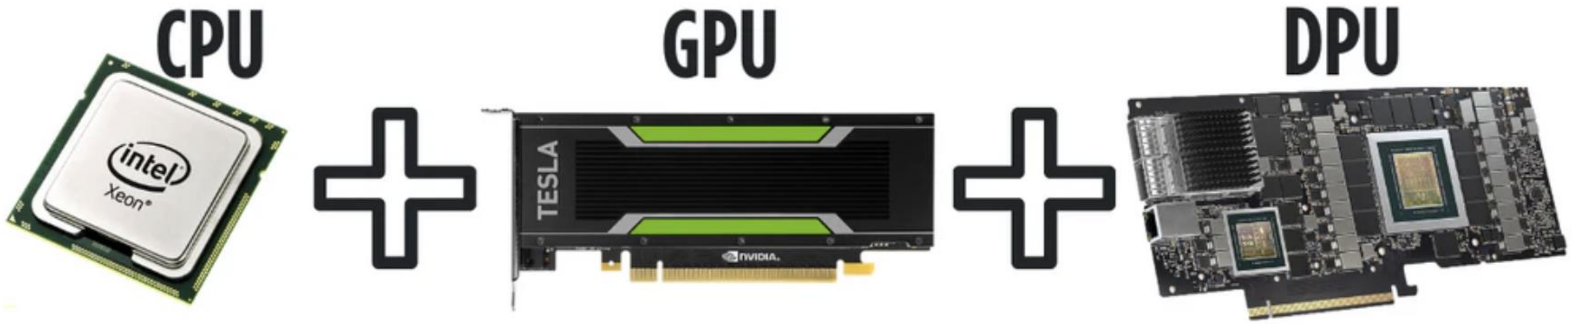

Source: Premionic

https://t.me/learningnets

and TPU

DPU vs GPU vs TPU:

Flexibility:

GPUs are the most flexible, as they can handle many tasks.

DPUs and TPUs are more specialized.

and TPU

DPU vs GPU vs TPU:

Performance:

TPUs excel in AI, GPUs in deep learning, but GPUs handle non-AI tasks too.

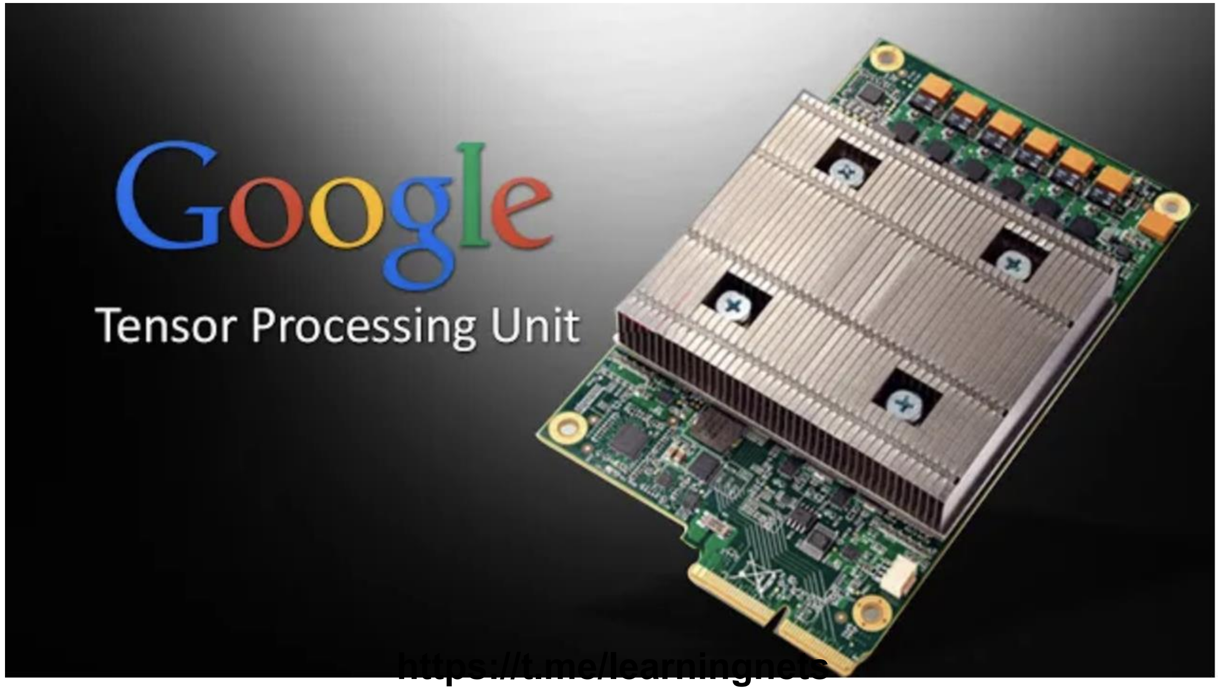

and TPU

DPU vs GPU vs TPU:

DPUs don’t do computations like GPUs or TPUs.

They offload non-computational tasks to improve infrastructure efficiency.

and TPU

DPU vs GPU vs TPU:

Is a DPU less powerful?

A DPU isn't about being more or less powerful than GPUs or TPUs.

It's built for different tasks, offloading and speeding up non-AI parallel computing jobs.

and TPU

DPU vs GPU vs TPU:

GPUs and TPUs are built for heavy computation.

GPUs handle parallel tasks like AI training, graphics, and simulations.

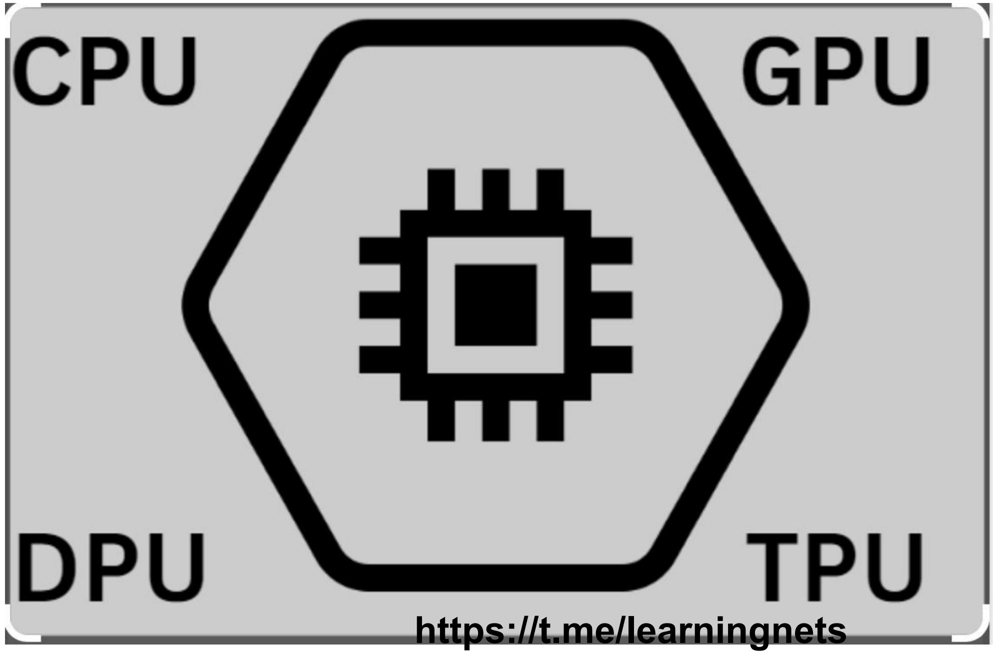

and TPU

DPU vs GPU vs TPU:

TPUs are built specifically for tensor-based AI tasks like neural networks.

DPUs handle networking, storage, and security, taking care of tasks like data transfer, encryption, and virtual machine management, freeing the CPU.

and TPU

DPU vs GPU vs TPU:

DPUs don’t handle heavy AI tasks like GPUs or TPUs but ensure the infrastructure runs smoothly.

They aren’t used to train large neural networks or process massive matrix operations.

and TPU

DPU vs GPU vs TPU:

Instead, DPU is used to handle data-centric tasks that would otherwise had to be done by CPU.

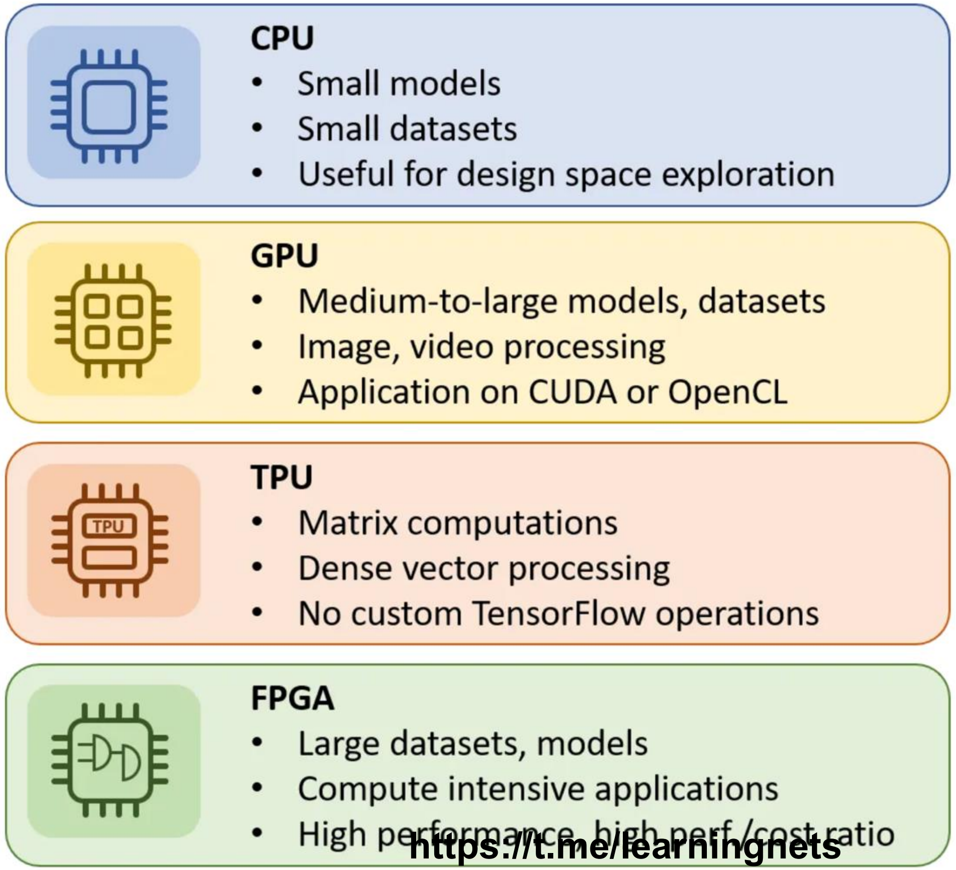

and TPU

DPU vs GPU vs TPU:

Why do we need DPUs? Can’t CPUs handle the same tasks?

No, CPUs are not designed for packet processing, encryption, decryption, or storage management.

These tasks are data-heavy and consume CPU cycles, which slows down applications.

and TPU

DPU vs GPU vs TPU:

Basically, DPUs are like a network ASICs.

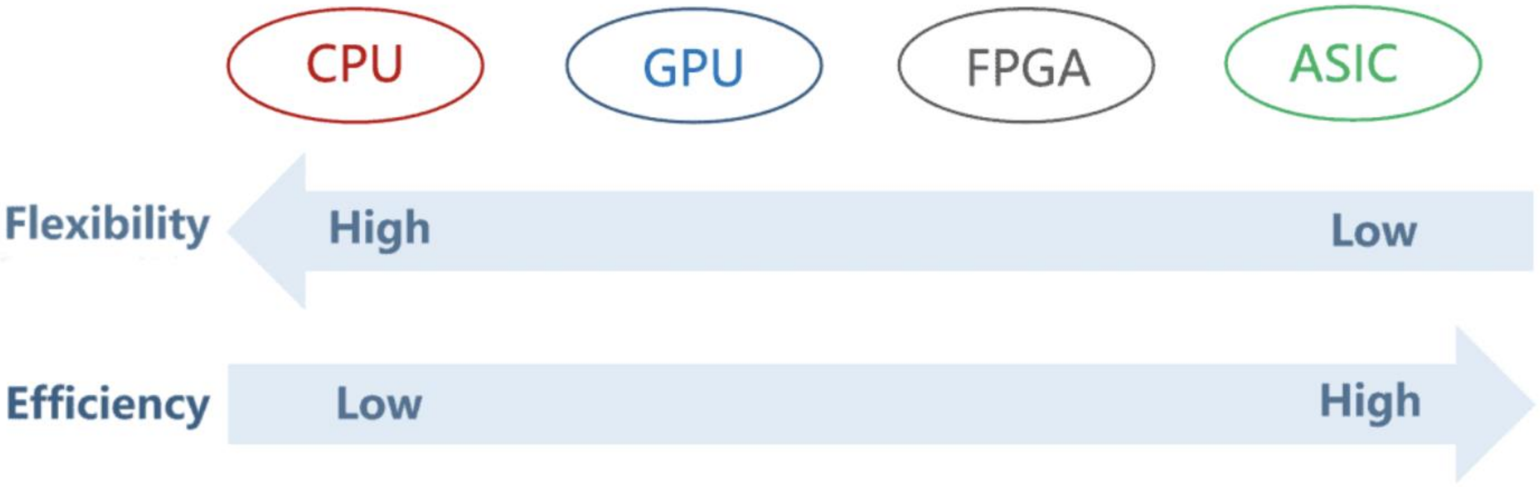

and TPU

Summary:

GPUs are the most general-purpose, DPUs are for data movement and resource utilization, and TPUs provide excellent efficiency for Google-based AI workloads.

agenda

NVMe over Fabrics

NVLink is NVIDIA’s high-speed interconnect technology.

It is for communication between GPUs, and sometimes between GPUs and CPUs **inside a single server.**

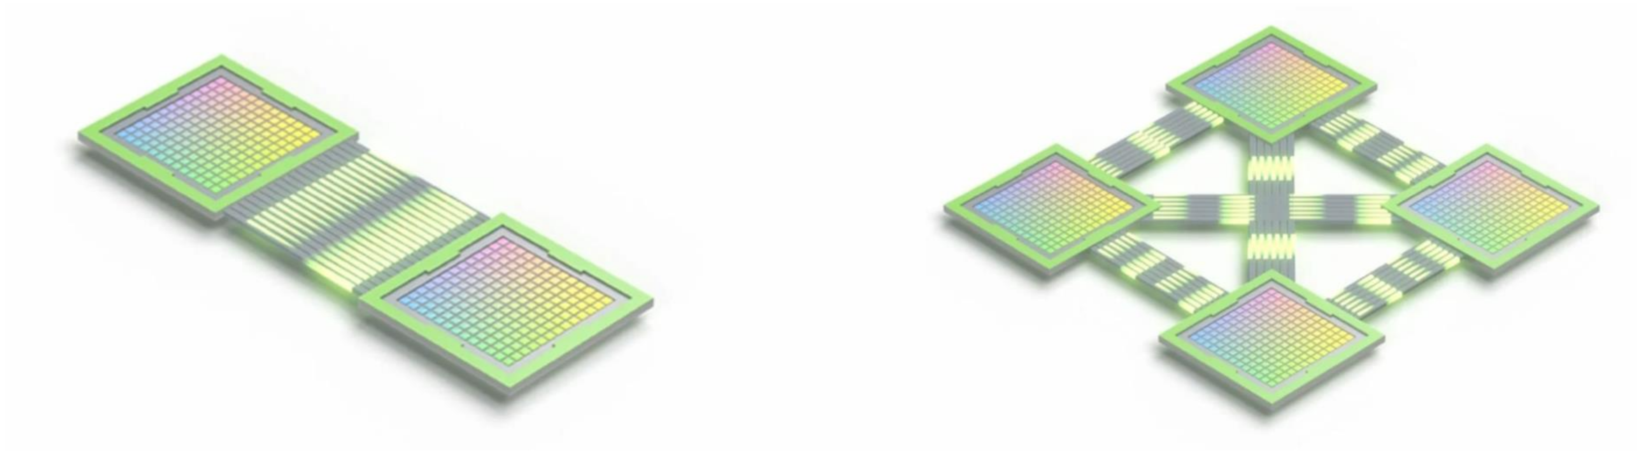

NVLink GPU to GPU Communication

NVMe over Fabrics

Infiniband and RoCE/RoCEv2 are used for inter-node communication; NVLink is just intra-node communication.

We can think of NVLink as a point-to-point physical communication link.

NVMe over Fabrics

With NVLink, GPUs can share memory directly similar to Infiniband and RoCE but as I said, scope is different.

NVLink is within a single server, between multiple GPUs in that server and the others are used to build entire fabric, inter-node communication

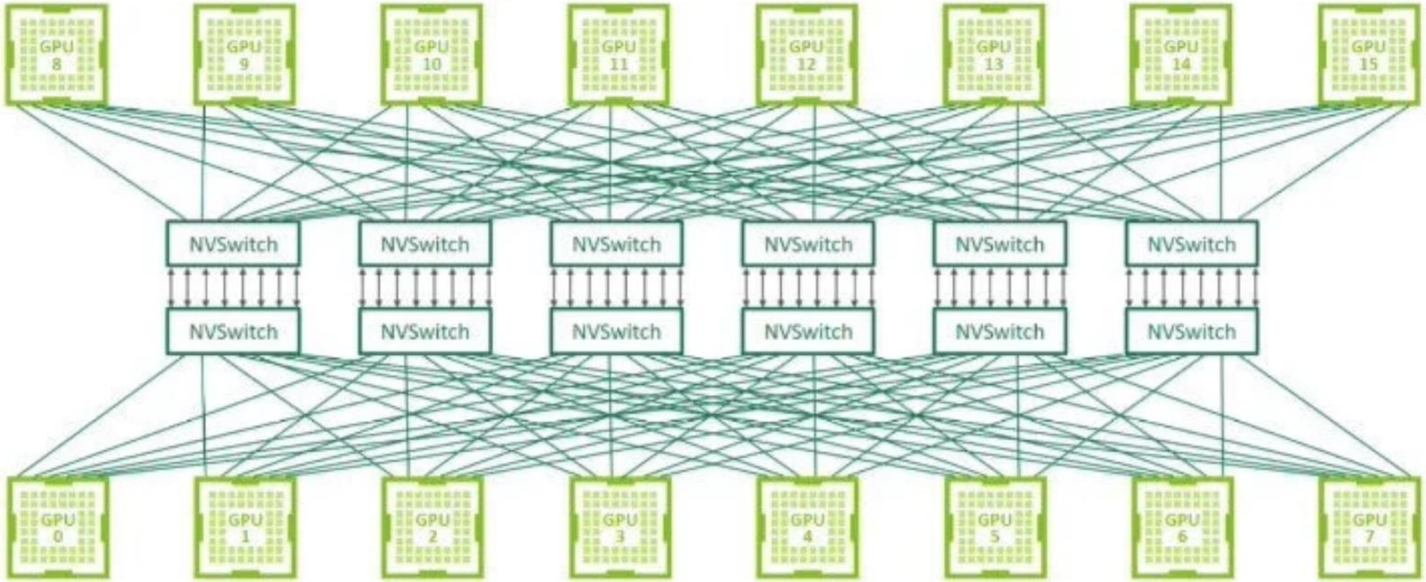

Nvdia NVswitch

NVMe over Fabrics

There is also NVSwitch, which is an extension of NVLink.

Think of it as a switching fabric for NVLink

NVMe over Fabrics

While NVLink connects GPUs directly in a peer-to-peer manner, NVSwitch provides communication between multiple GPUs in a system.

NVSwitch and NVLink together can be compared to Infiniband or ROCE

NVMe over Fabrics

NVLink, as of 2025, provides up to 900 GB/s of bandwidth per GPU in systems like NVIDIA's H100.

NVMe over Fabrics

NVLink doesn't involve Ethernet or RDMA; it is a proprietary NVIDIA interconnect for GPU communication.

NVMe over Fabrics

RoCEv2 and Infiniband can work across IP networks using ECN, DCQCN, or credit-based methods for congestion control.

NVLink, on the other hand, is strictly used for GPU-to-GPU communication in a single node.

NVMe over Fabrics

NVMe over Fabrics is used to extend NVMe protocol across a network.

Accessing NVMe storage devices remotely with minimal latency, providing near-local performance for distributed storage systems.

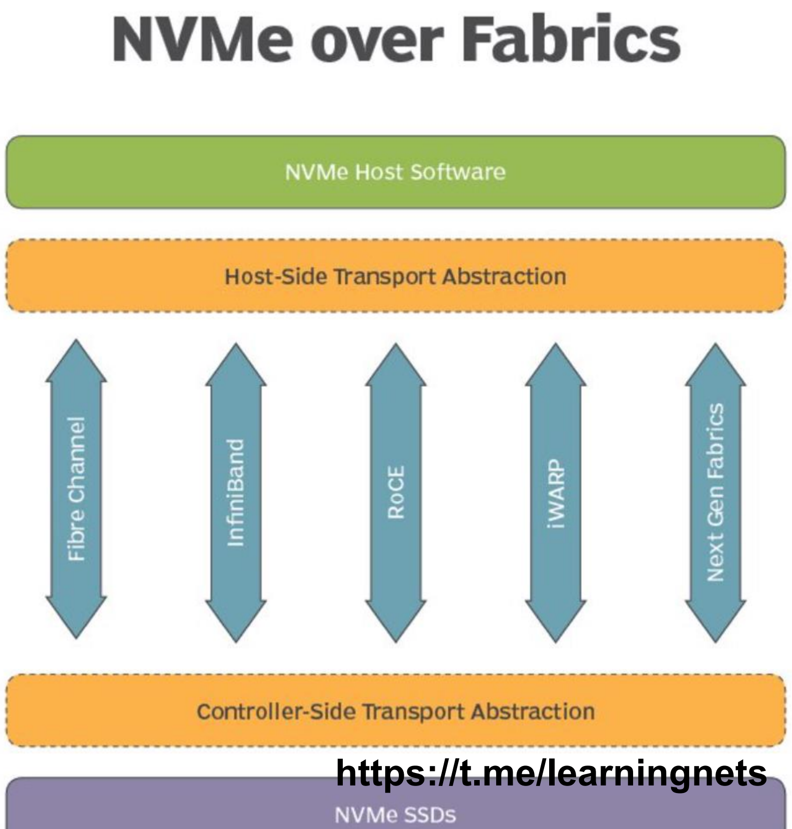

NVMe over Fabric

In this design, storage can be centralized and shared between many computing resources, GPUs.

Traditional NVMe is designed for direct-attached storage (DAS) , within a single system.

NVMe over Fabric

NVMe over Fabric can use Ethernet, Infiniband, and Fiberchannel as an underlying transport; it is a transport agnostic.

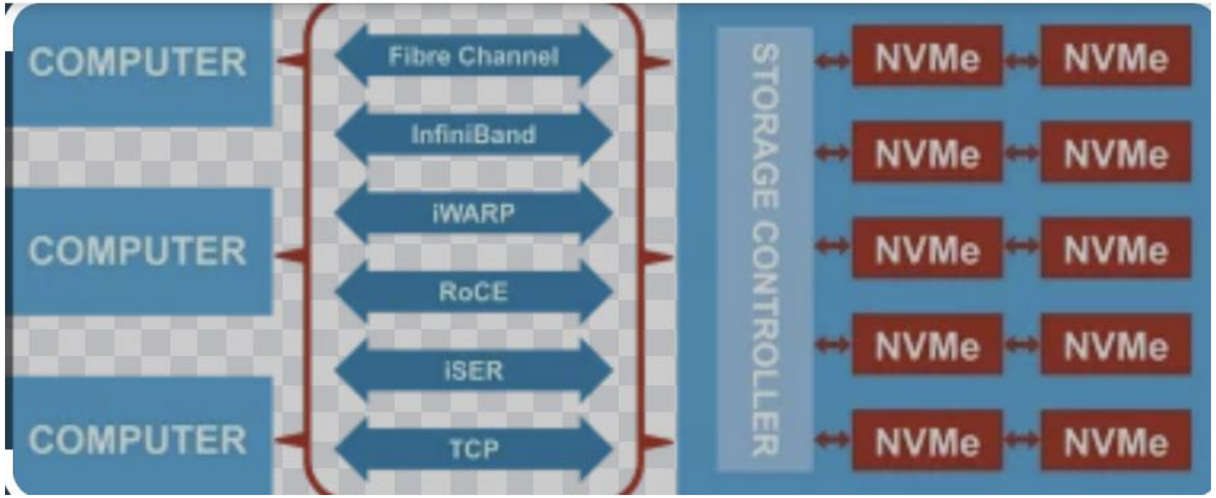

DCBX

In real networks, we see both Layer 2 and Layer 3 protocols are used together.

For example, PFC and ECN are used together in many networks.

DCBX

When it comes to layer 2-based congestion control protocols, not just one protocol but a suite of protocols need to be discussed.

That suite of protocols is called DCB, or Datacenter Bridging.

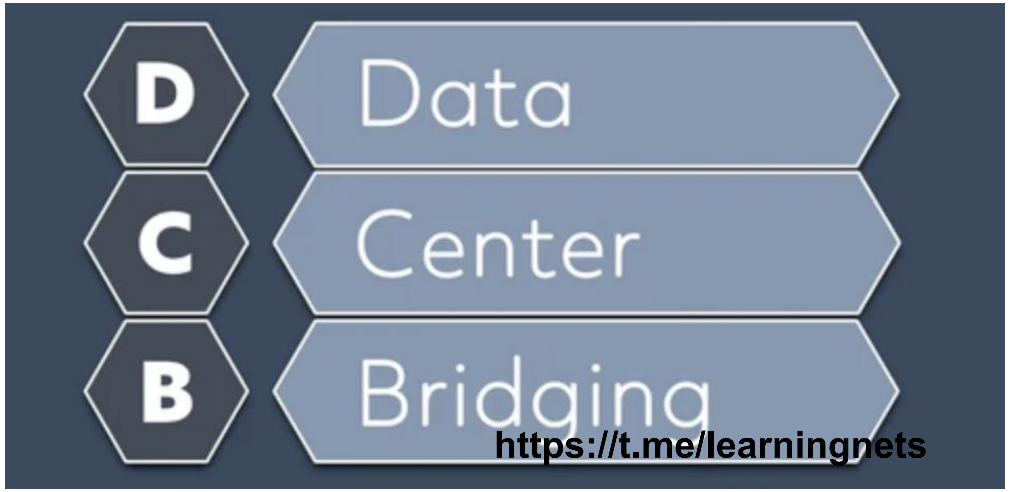

DCBX

QCN, Quantized Congestion Notification, DCBX, Datacenter Bridging Exchange, ETS, Enhanced Tranmission Selection and last but not least PFC, Priority Flow Control.

All of them are DCB protocols.

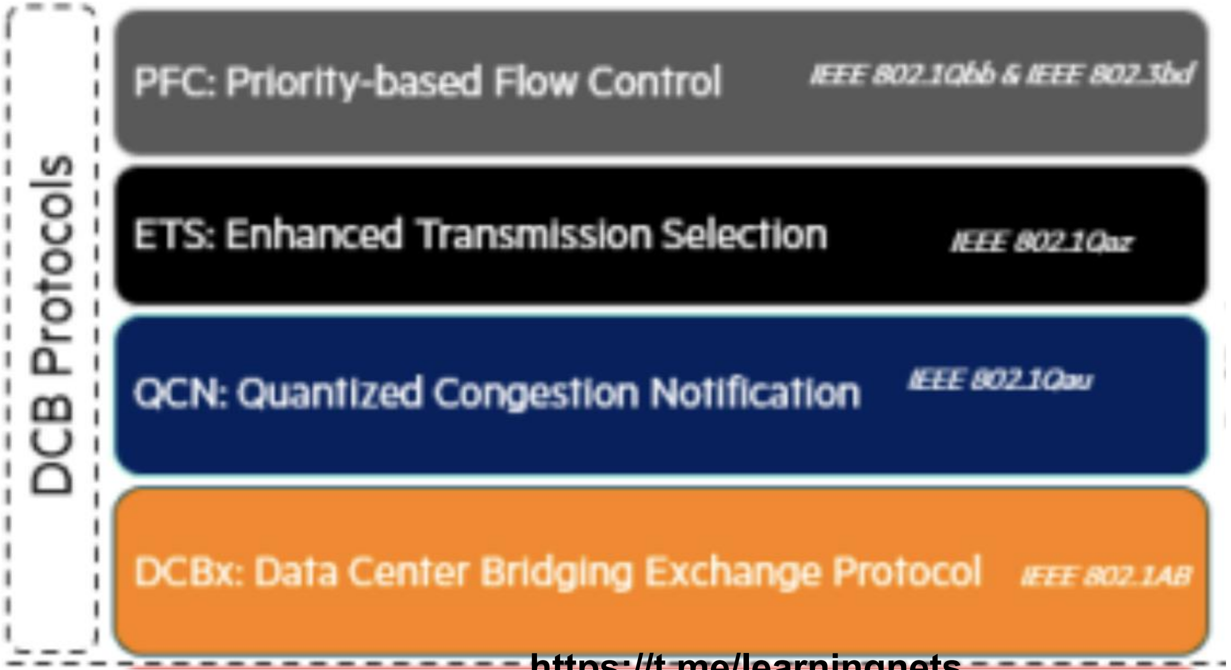

Source: Arubanetworking

DCBX

DCBX is a discovery and dynamic DCB parameter negotiation protocol basically.

It allows devices like switches or servers to exchange and negotiate DCB settings dynamically.

Similar to LACP, DTP, etc., dynamic negotiation.

Devices advertise their supported DCB features, PFC, QCN and ETS and negotiate their parameters with that.

ETS – Enhanced Transmission Selection

Enhanced Transmission Selection which is IEEE 802.1Qaz standard.

ETS allocates bandwidth to different traffic classes, so critical traffic like RDMA gets minimum bandwidth guarantee in AI.

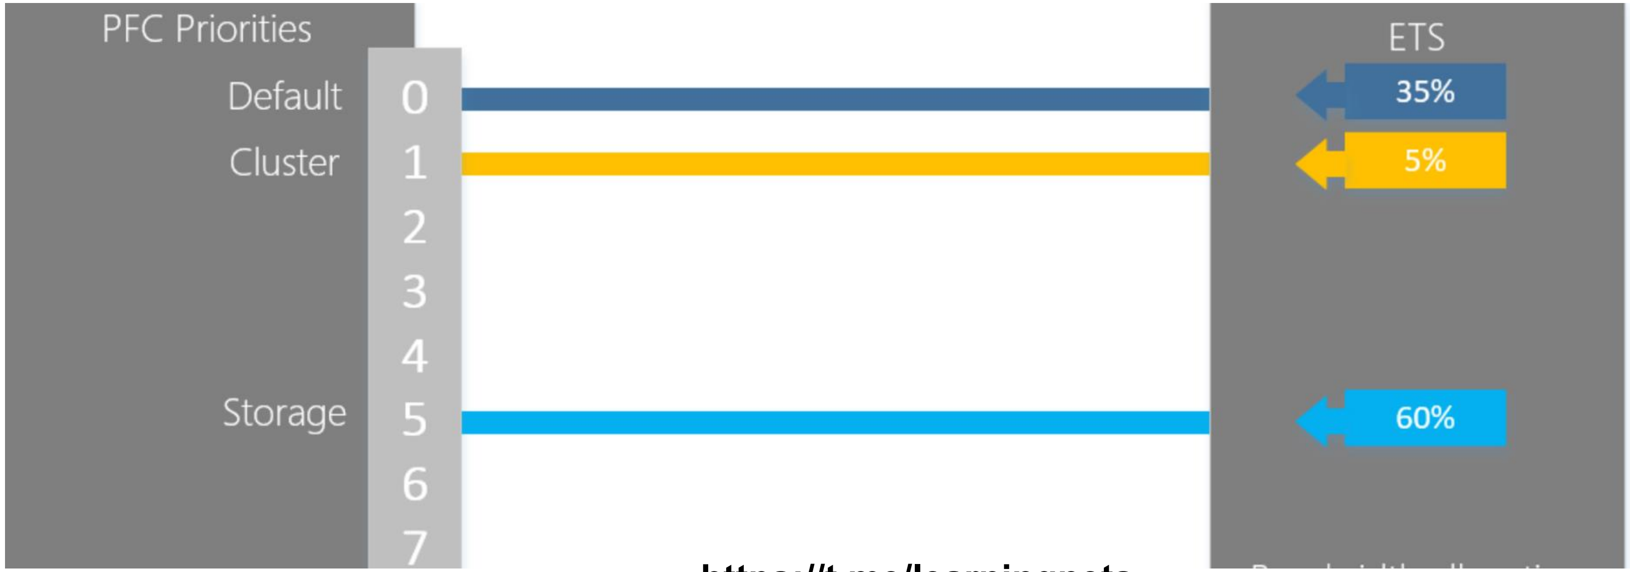

ETS – Enhanced Transmission Selection

Traffic classes are assigned percentages of link bandwidth, so important traffic gets a minimum bandwidth guarantee.

ETS is similar to RSVP so it provides explicit bandwidth allocation, but it’s limited to individual Ethernet links, and operates on traffic groups rather than individual flows, like in RSVP.

Allocation bandwidth, basically but in a link level, not end-to-end like RSVP.

PFC – Priority Flow Control

Priority Flow Control which is IEEE 802.1Qbb standard.

Probably the most important and debated technology in Congestion Control is PFC.

Because when it is used, certain traffic is dropped and it can cause other problems such as Head of Line Blocking etc.

PFC – Priority Flow Control

PFC is a networking protocol that allows devices to pause and resume the transmission of data based on the priority of the traffic.

It is used to prevent congestion and ensure that high-priority traffic is not delayed by lower-priority traffic.

PFC – Priority Flow Control

With PFC, devices send special pause frames to their neighbors, indicating that they are temporarily unable to receive more data.

The neighbors will then stop sending data until they receive a resume signal.

PFC – Priority Flow Control

PFC provides lossless Ethernet by pausing traffic for specific priority classes when congestion happens.

PFC – Priority Flow Control

Ethernet frames are tagged with priority levels, 0 to 7, and PFC can pause traffic for selected priorities while allowing others to continue.

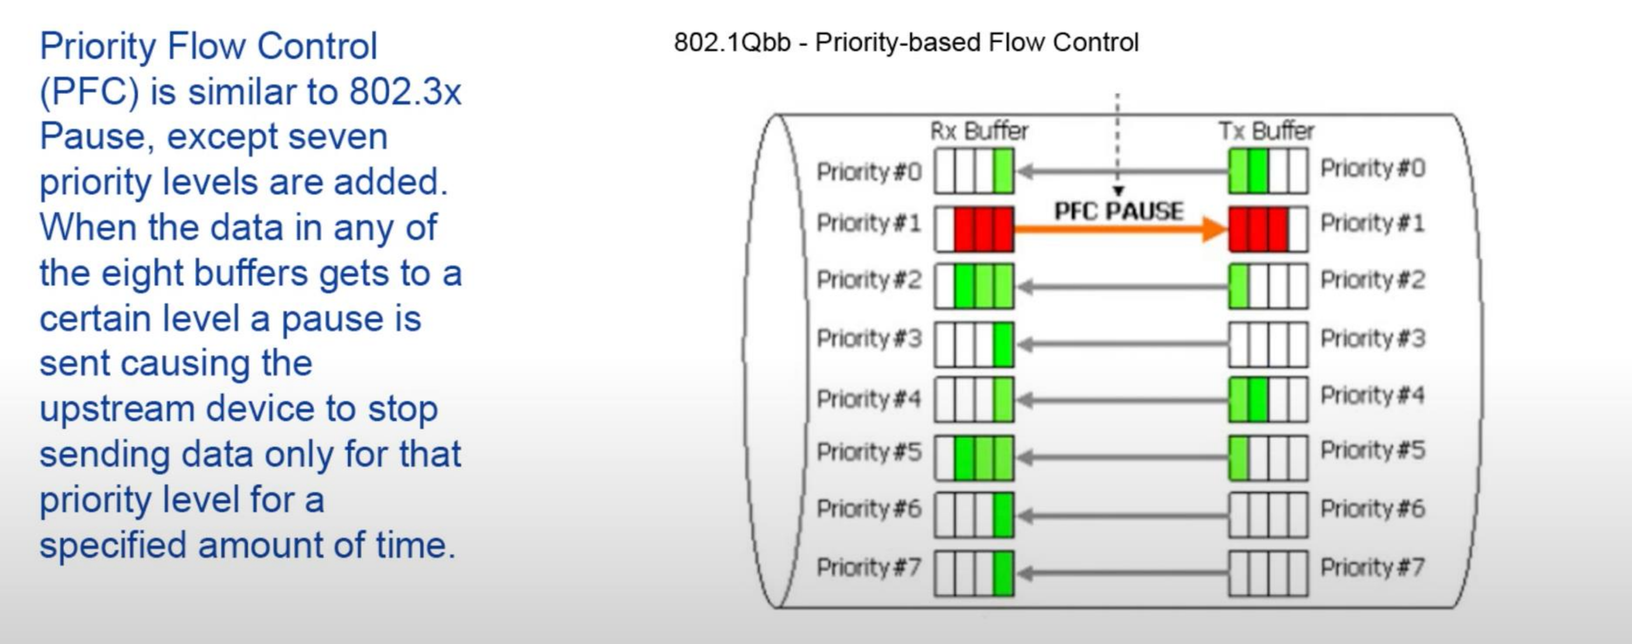

Source: Flash Memory Summit

PFC – Priority Flow Control

For example, RDMA, for the AI applications, should continue, while others stop.

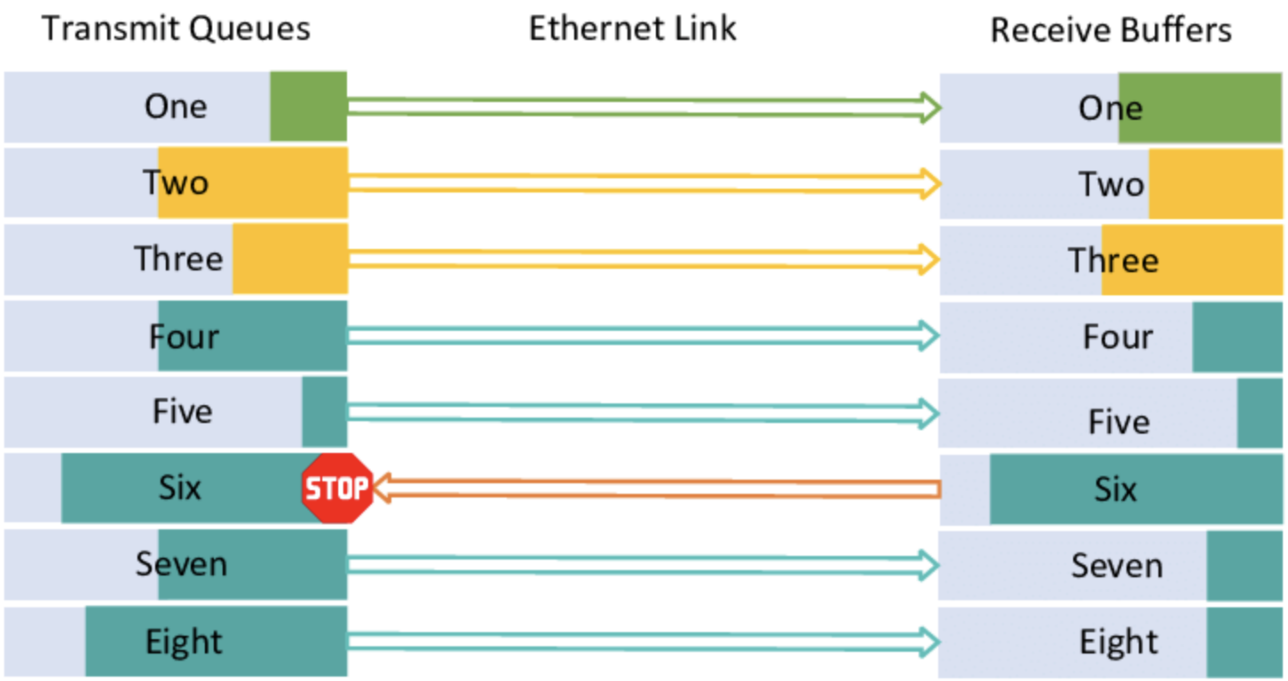

Source: aboutnetworks.net

QCN – Quantized Congestion Notification

QCN signals the sender to adjust sending rate when congestion is detected.

In QCN, Switches monitor congestion and send Congestion Notification Messages, CNMs back to the sender,and sender adjusts the sending rate basically.

QCN – Quantized Congestion Notification

QCN is not the same as DCQCN.

QCN is part of the DCB, which are all layer 2 mechanisms

DCQCN is Layer 3, an end-to-end congestion control mechanism.

QCN – Quantized Congestion Notification

Summary:

DCBX : Exchanging parameters
PFC : Pausing traffic
ETS: Allocating bandwidth for some traffic
QCN: Notification about the congestion

All are layer 2 and part of DCB, Datacenter Bridging.

Protocol	Purpose	Key Use Case
PFC (802.1Qbb)	Provides lossless Ethernet by pausing traffic for priorities	Used in RoCE and FCoE for reliable delivery.
ETS (802.1Qaz)	Allocates bandwidth to traffic classes	Balances bandwidth for RDMA, storage, and Ethernet traffic
QCN (802.1Qau)	Handles congestion via feedback in Layer 2 networks	Prevents congestion-related packet drops in Ethernet
DCBX	Dynamically exchanges and configures DCB settings	Ensures consistent DCB configurations across devices.

ECN

There are many Layer 3 congestion control protocols.

I will cover ECN-based and RTT-based congestion control protocols in this course.

ECN, Explicit Congestion Notification, DCQCN, Data Center Quantized Congestion Notification, TIMELY, and HPCC, High Precision Congestion Control protocol will be explained.

ECN

These technologies are used to manage and optimize network traffic for AI workloads

ECN

ECN, Explicit Congestion Notification, in TCP/IP networks to signal congestion without dropping packets.

Normally, in traditional congestion control, when a network experiences congestion, packets get dropped. RED/WRED etc.

This packet loss acts as a signal to the sender to slow down.

ECN

Packet drops can be inefficient, especially in high-performance environments, because due to retransmissions, latency increases too.

ECN solves this by marking packets to indicate congestion instead of dropping them.

ECN

When a switch or router in the network detects that its buffers are getting full, it marks the packets with specific bits in the IP header.

This marking tells the receiver that the path is congested.

ECN

When the receiver gets an ECN-marked packet, it notifies the sender about the congestion through an acknowledgment message.

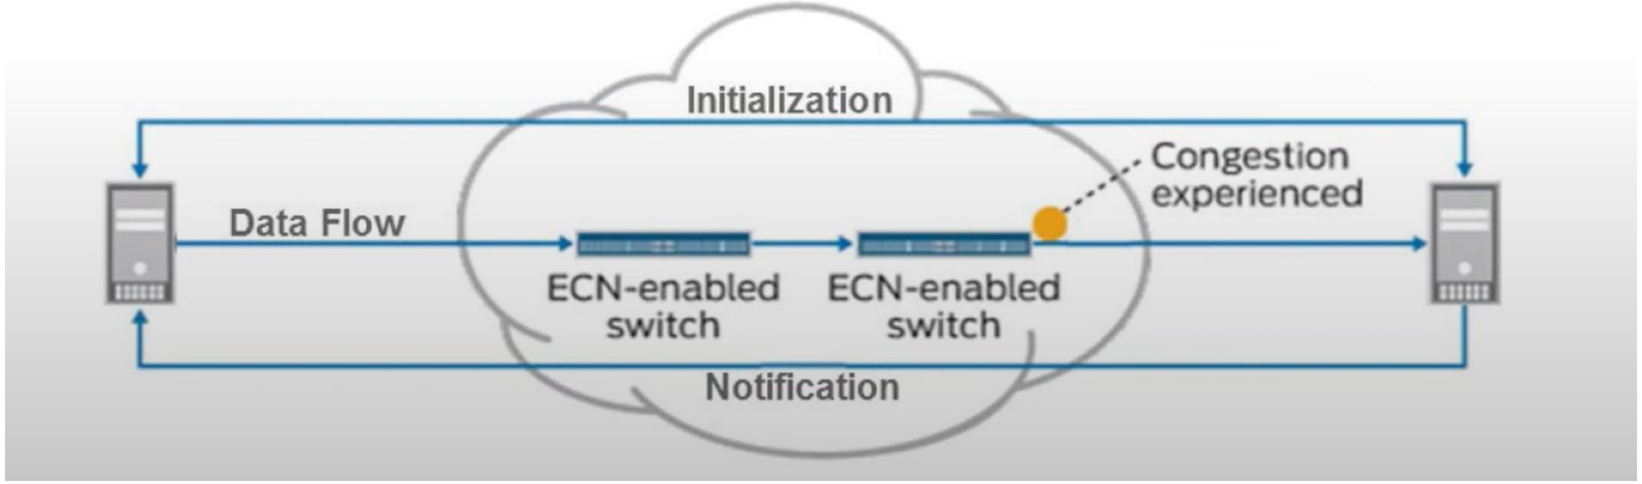

ECN Notification

ECN

Switches notify the receiver, which then informs the sender, adding extra latency.

Later in this video, I'll explain how protocols like Timely Protocol avoid this by removing the receiver from the process, making it faster.

agenda

ECN

When the sender gets the signal from the receiver, it slows down to manage congestion.

ECN helps the network handle congestion early, preventing packet loss.

Before buffers overflow, ECN lets end hosts communicate and adjust their speeds without dropping packets.

DCQCN

DCQCN relies on ECN to provide congestion control specifically for RoCEv2 traffic, but not with RoCE

While ECN signals congestion, DCQCN improves it by introducing a more granular way to manage congestion

DCQCN

ECN only notifies the sender that congestion is happening, but it doesn't give details about how severe the congestion is!

DCQCN uses ECN markings as the foundation for its congestion notification.

DCQCN

How does DCQCN work?

With normal ECN, the receiver sends ACK for notification but with DCQCN, the receiver generates a special feedback packet called a Congestion Notification Packet (CNP).

CNP contains more detailed, quantized feedback about the congestion level, allowing the sender to adjust its sending rate precisely.

DCQCN

When the sender receives a CNP, it adjusts its sending rate based on the severity of congestion.

If the congestion is light, the sender reduces its rate slightly.

If the congestion is severe, the sender decreases the rate more significantly.

DCQCN

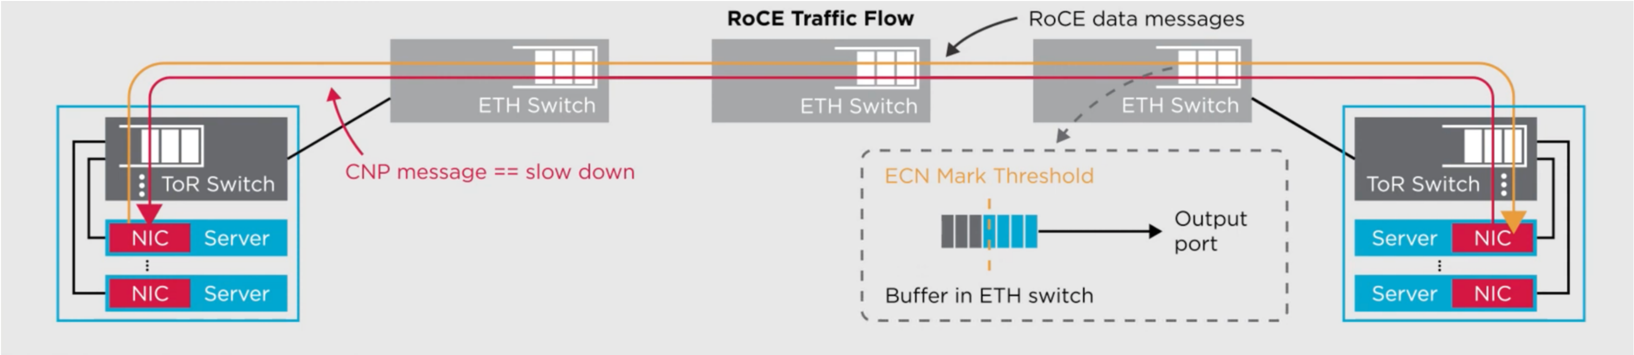

CNP – Congestion Notification Packet

https://t.me/learningnets

DCQCN vs ECN

ECN alone is binary, so it can only indicate whether congestion exists or not.

It doesn't provide granular feedback on the severity of congestion.

DCQCN does, through quantized feedback via CNP!

DCQCN

DCQCN vs ECN

When congestion stops, the sender gradually increases its transmission rate.

By improving ECN with detailed, quantized feedback, DCQCN provides much more precise congestion control than normal ECN.

Feature	Normal ECN	Quantized Feedback (DCQCN)
Congestion Signal	Binary (Marked or Not Marked)	Granular (Quantized Level of Congestion)
Reaction by Sender	Reactive, often abrupt	Proactive, gradual, and measured
Rate Adjustment	Fixed or arbitrary adjustments	Dynamically calculated based on severity
Scalability	Limited in multi-flow environments	Scales well for large-scale RDMA flows

DCQCN vs ECN

ECN and DCQCN have different ways of signaling and handling congestion.

CNP is sent in DCQCN to indicate the severity of the congestion.

But ECN itself does not send CNPs. It sends ACK, meaning there is congestion, no ACK means no congestion.

CNPs are part of the DCQCN only!

ECN More Detail:

In ECN, the switch or router monitors its buffer occupancy.

If the buffer exceeds the threshold, it declares congestion.

Switch marks the CE and Congestion Experienced bits in the IP header

DCQCN

ECN More Detail:

The CE bits are part of the ECN field in the IP header

00: Non-ECN capable, 01 or 10: ECN capable, and 11: Congestion Experienced, CE.

ECN More Detail:

If the CE bit is set by the switch, the receiver understands that the path is congested.

The receiver then notifies the sender of the congestion using a flag in the TCP ACK.

Specifically, the ECE, ECN Echo bit in the TCP header is set to tell the sender that the network is congested.

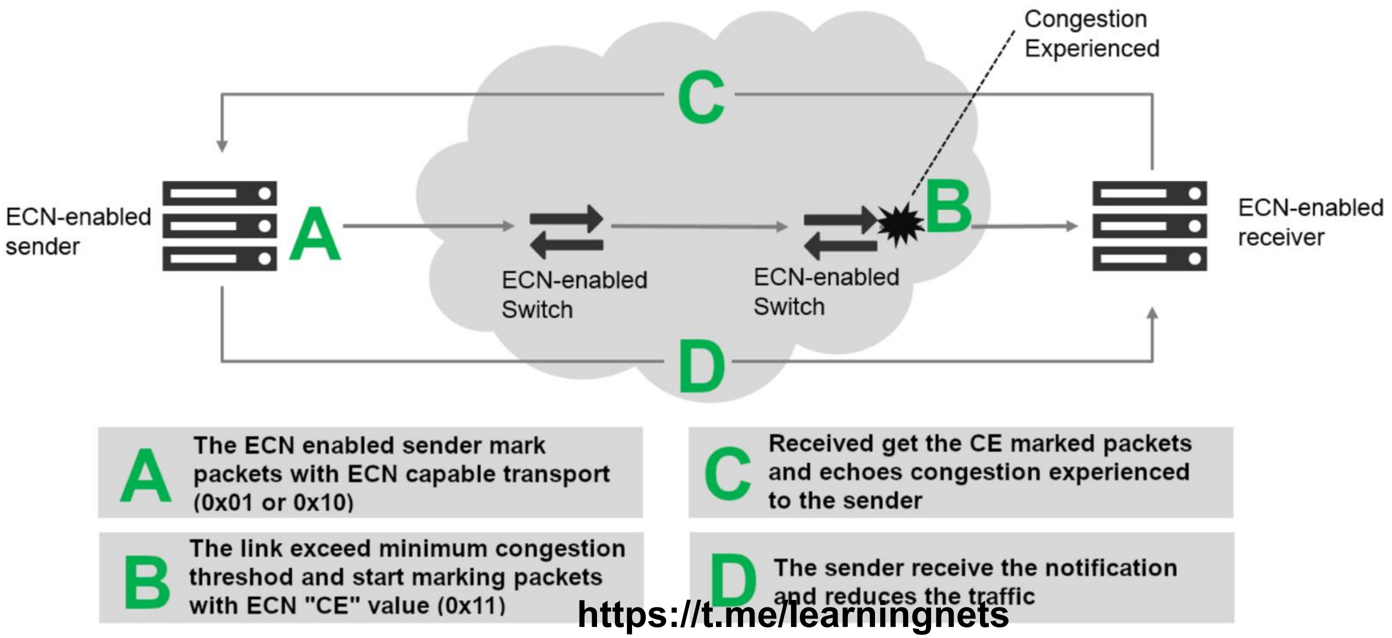

ECN More Detail:

When the sender receives an ACK with the ECE bit, it understands that there is congestion.

The sender reduces its transmission rate.

The sender, after reducing its rate, sets the CWR, Congestion Window Reduced bit

ECN More Detail:

When the receiver sees the CWR bit, it stops setting the ECE bit in further ACKs!

In DCQCN CE bit is sent to the receiver, but instead of basic TCP ACK with the ECE bit, the receiver sends CNP, Congestion Notification Packet to the sender as quantized feedback (severity of congestion)

ECN vs DCQCN

One of the important differences between ECN and DCQCN is, ECN works for Layer 3 TCP/IP traffic in general networks, but DCQCN is specifically designed for RDMA traffic over RoCEv2, a specialized Layer 3 transport.

DCQCN

DCQCN is a combination of ECN and PFC

In DCQCN, ECN is used to address the drawbacks of PFC, enabling a lossless Ethernet network.

The concept of DCQCN is to let ECN handle flow control by reducing the sending rate when congestion begins, which helps delay or minimize the activation of PFC, which completely pauses traffic.

DCQCN

The goal of DCQCN is to balance two opposing needs:

1: Make sure PFC doesn't activate too soon, giving ECN enough time to provide congestion feedback and slow down the traffic.

2: Make sure PFC doesn't activate too late, as this could lead to packet loss from buffer overflow.

DCQCN

In short, PFC should trigger at the right moment, not too early, not too late.

If PFC is necessary, use it, otherwise, let ECN handle congestion.

DCQCN

QCN vs DCQCN

Both QCN and similar systems work to avoid packet loss by alerting the sender about congestion and allowing it to adjust its transmission rate dynamically.

But, QCN operates as a Layer 2 congestion notification method.

QCN vs DCQCN

Feature	QCN (Layer 2)	DCQCN (Layer 3+ for RoCEv2)
Layer	Operates at Layer 2 (Ethernet)	Operates at Layer 3 (IP) with ECN
Congestion Detection	Switch monitors buffer utilization	Switch marks packets with ECN bits
Congestion Feedback	Switch sends a Congestion Notification Message (CNM) to the sender	Receiver sends a Congestion Notification Packet (CNP) to the sender
Granularity	Traffic-class level (no per-flow control)	Per-flow level (granular control)
Scope	Limited to Laver 2 domains (single Ethernet network)	Scales across large, multi-hop Layer 3 IP networks)
Use Case	Used in RoCE (Laver 2) and FCoE networks	Used in RoCEv2 (Laver 3) for RDMA in large-scale data centers
Implementation Complexity	Simpler to implement, limited to Layer 2	More complex, requires ECN and CNP support

QCN vs DCQCN Comparison Chart

Timely

Timely is a VMware protocol for RoCEv2 networks.

It’s gaining attention for its importance in modern networking.

Timely enhances RDMA performance over Ethernet in RoCEv2.

Timely

It uses delay-based control, detecting congestion via round-trip time (RTT).

The goal is to adjust transmission rates based on delay to prevent congestion.

Timely

It measures the RTT between sender and receiver to detect congestion.

Timely identifies congestion by tracking packet delays (RTT).

An increase in RTT signals growing congestion.

Timely

It adjusts the sending rate dynamically based on this delay.

Timely aims to prevent congestion by proactively managing traffic rates, similar to DCQCN and ECN.

Timely

Unlike DCQCN, which uses ECN to mark packets, Timely relies on RTT (latency) for congestion detection.

Feature	Timely	DCQCN
Congestion Detection	Based on RTT (latency measurements)	Based on ECN (Explicit Congestion Notification)
Feedback Sourse	Sender-side only (latency-based feedback)	Switches mark ECN bits; receiver sends feedback to the sender)
Granularity	Per-flow congestion control	Per-flow congestion control
Reaction Speed	Faster reaction (real-time RTT-based)	Slightly slower (relies on ECN marking and receiver feedback)
Scalability	Suitable for smaller or mid-sized networks	Scales better in large networks with high traffic
Hardware Dependency	Requires no specific switch support for ECN	Requires ECN-capable switches and devices
Accuracy	Less precise in detecting actual congestion sources	More precise due to switch-based congestion marking
Complexity	Simpler to implement	More complex depends on switch and endpoint configurations)

Timely vs DCQCN Comparison Chart

Timely

Latency acts as an early signal, helping avoid full congestion.

Timely adjusts rates based on delay, not ECN markings or telemetry, unlike DCQCN and HPCC.

HPCC – High Precision Congestion Control

HPCC is a newer protocol designed for AI training clusters.

It improves precision and responsiveness compared to other Layer 3 congestion control protocols.

HPCC – High Precision Congestion Control

HPCC relies on in-network telemetry to detect real-time congestion and reacts faster than DCQCN.

Telemetry data helps the sender adjust its rate with greater speed and accuracy.

HPCC – High Precision Congestion Control

Like DCQCN and Timely, HPCC is built for large-scale RDMA-based data centers, especially those using RoCEv2.

It uses telemetry from switches to identify congestion.

.

Feature	Timely	HPCC
Congestion Detection	Based on RTT (delay inference)	Based on in-network telemetry (queue depth)
Feedback Source	End-host RTT measurement	Real-time telemetry from switches
Precision	Moderate (delay-based, less granular)	High (telemetry-driven, very granular)
Responsiveness	Slower and smoother	Faster and more accurate
Scalability	Moderate (for small to medium deployments)	High (suitable for large-scale data centers)
Implementation Complexity	Simple (end-host only)	Complex (requires telemetry-capable switches)
Use Case	Moderate-scale RoCEv2 deployments	Large-scale AI clusters and HPC workloads

Timely vs HPCC Comparison Chart

HPCC – High Precision Congestion Control

HPCC uses real-time telemetry, like queue depth and traffic volume, instead of delay or buffer usage.

This provides precise congestion feedback to the sender.

.

HPCC – High Precision Congestion Control

Not everyone uses HPCC because it requires costly, telemetry-capable switches.

While complex to implement, it offers high accuracy and control.

.