AI & Blockchain

Module Quiz – 20 Multiple-Choice Questions

Topics covered: Convergence thesis · Decentralized compute · Tokenized AI · zkML · AI agents · Content provenance · Risks · Evaluation framework

Select the best answer. Answers revealed after each question.

Bloom's levels: 4 Understand · 8 Apply · 6 Analyze · 2 Evaluate

---

20 questions covering the full AI & Blockchain lecture. Time: approximately 15 minutes.

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?
A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale
C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?

A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale

C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks

Answer: (C) – The thesis holds that AI and blockchain solve different problems and reinforce each other: AI adds intelligence, blockchain adds verifiability and decentralized coordination.

Q2 [Understand]. What coordination problem does decentralized compute aim to solve?

A) The lack of internet connectivity in developing countries B) Concentration of GPU capacity among a few hyperscalers creating access barriers and price power

C) The high energy consumption of proof-of-work blockchains D) The inability of open-source models to match proprietary model quality

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?
A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale
C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks
Answer: (C) – The thesis holds that AI and blockchain solve different problems and reinforce each other: AI adds intelligence, blockchain adds verifiability and decentralized coordination.

Q2 [Understand]. What coordination problem does decentralized compute aim to solve?
A) The lack of internet connectivity in developing countries B) Concentration of GPU capacity among a few hyperscalers creating access barriers and price power
C) The high energy consumption of proof-of-work blockchains D) The inability of open-source models to match proprietary model quality
Answer: (B) – Decentralized compute networks (Akash, Render, io.net) aggregate idle GPU supply globally to reduce dependence on AWS, Azure, and GCP.

Q3 [Understand]. What is Zero-Knowledge Machine Learning (zkML)?
A) A method for training AI models without labeled data B) A technique for compressing large language models
C) A cryptographic method proving a model produced a specific output without revealing the model weights or input D) A blockchain protocol for storing model parameters on-chain

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?
A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale
C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks
Answer: (C) – The thesis holds that AI and blockchain solve different problems and reinforce each other: AI adds intelligence, blockchain adds verifiability and decentralized coordination.

Q2 [Understand]. What coordination problem does decentralized compute aim to solve?
A) The lack of internet connectivity in developing countries B) Concentration of GPU capacity among a few hyperscalers creating access barriers and price power
C) The high energy consumption of proof-of-work blockchains D) The inability of open-source models to match proprietary model quality
Answer: (B) – Decentralized compute networks (Akash, Render, io.net) aggregate idle GPU supply globally to reduce dependence on AWS, Azure, and GCP.

Q3 [Understand]. What is Zero-Knowledge Machine Learning (zkML)?
A) A method for training AI models without labeled data B) A technique for compressing large language models
C) A cryptographic method proving a model produced a specific output without revealing the model weights or input D) A blockchain protocol for storing model parameters on-chain
Answer: (C) – zkML uses zero-knowledge proofs to prove model execution integrity – enabling verifiable inference while preserving privacy of inputs and proprietary model parameters.

Q4 [Understand]. What distinguishes an on-chain AI agent from a standard smart contract?
A) On-chain AI agents are written in Solidity; smart contracts use Python B) Smart contracts execute deterministic logic; AI agents can reason, plan, and take multi-step actions autonomously
C) AI agents operate on Layer 2 only; smart contracts run on Layer 1 D) AI agents replace oracles by reading blockchain state directly

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?
A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale
C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks
Answer: (C) – The thesis holds that AI and blockchain solve different problems and reinforce each other: AI adds intelligence, blockchain adds verifiability and decentralized coordination.

Q2 [Understand]. What coordination problem does decentralized compute aim to solve?
A) The lack of internet connectivity in developing countries B) Concentration of GPU capacity among a few hyperscalers creating access barriers and price power
C) The high energy consumption of proof-of-work blockchains D) The inability of open-source models to match proprietary model quality
Answer: (B) – Decentralized compute networks (Akash, Render, io.net) aggregate idle GPU supply globally to reduce dependence on AWS, Azure, and GCP.

Q3 [Understand]. What is Zero-Knowledge Machine Learning (zkML)?
A) A method for training AI models without labeled data B) A technique for compressing large language models
C) A cryptographic method proving a model produced a specific output without revealing the model weights or input D) A blockchain protocol for storing model parameters on-chain
Answer: (C) – zkML uses zero-knowledge proofs to prove model execution integrity – enabling verifiable inference while preserving privacy of inputs and proprietary model parameters.

Q4 [Understand]. What distinguishes an on-chain AI agent from a standard smart contract?
A) On-chain AI agents are written in Solidity; smart contracts use Python B) Smart contracts execute deterministic logic; AI agents can reason, plan, and take multi-step actions autonomously
C) AI agents operate on Layer 2 only; smart contracts run on Layer 1 D) AI agents replace oracles by reading blockchain state directly
Answer: (B) – Smart contracts are deterministic and rule-based; AI agents add planning, language understanding, and adaptive decision-making on top of blockchain execution.

Q5 [Apply]. AWS charges $32/hr per GPU; Akash charges $5/hr. A training run requires 200 GPU-hours. What is the cost saving using Akash?
A) $1,400 B) $5,400 C) $640 D) $1,000

Quiz: Questions 1–5

Understand / Apply

Q1 [Understand]. What does the "Convergence Thesis" claim about AI and blockchain?
A) AI and blockchain are competing technologies that cannot coexist B) AI requires blockchain to function at scale
C) AI and blockchain are complementary: AI provides intelligence, blockchain provides trust and auditability D) Blockchain models replace AI models for prediction tasks
Answer: (C) – The thesis holds that AI and blockchain solve different problems and reinforce each other: AI adds intelligence, blockchain adds verifiability and decentralized coordination.

Q2 [Understand]. What coordination problem does decentralized compute aim to solve?
A) The lack of internet connectivity in developing countries B) Concentration of GPU capacity among a few hyperscalers creating access barriers and price power
C) The high energy consumption of proof-of-work blockchains D) The inability of open-source models to match proprietary model quality
Answer: (B) – Decentralized compute networks (Akash, Render, io.net) aggregate idle GPU supply globally to reduce dependence on AWS, Azure, and GCP.

Q3 [Understand]. What is Zero-Knowledge Machine Learning (zkML)?
A) A method for training AI models without labeled data B) A technique for compressing large language models
C) A cryptographic method proving a model produced a specific output without revealing the model weights or input D) A blockchain protocol for storing model parameters on-chain
Answer: (C) – zkML uses zero-knowledge proofs to prove model execution integrity – enabling verifiable inference while preserving privacy of inputs and proprietary model parameters.

Q4 [Understand]. What distinguishes an on-chain AI agent from a standard smart contract?
A) On-chain AI agents are written in Solidity; smart contracts use Python B) Smart contracts execute deterministic logic; AI agents can reason, plan, and take multi-step actions autonomously
C) AI agents operate on Layer 2 only; smart contracts run on Layer 1 D) AI agents replace oracles by reading blockchain state directly
Answer: (B) – Smart contracts are deterministic and rule-based; AI agents add planning, language understanding, and adaptive decision-making on top of blockchain execution.

Q5 [Apply]. AWS charges \$32/hr per GPU; Akash charges \$5/hr. A training run requires 200 GPU-hours. What is the cost saving using Akash?
A) \$1,400 B) \$5,400 C) \$640 D) \$1,000
Answer: (B) – AWS cost: $200 \times \$32 = \$6{,}400$. Akash cost: $200 \times \$5 = \$1{,}000$. Saving: $\$6{,}400 - \$1{,}000 = \$5{,}400$.

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs \$28/hr; a decentralized network costs \$6/hr. How much does the startup save?
A) \$11,000 B) \$3,000 C) \$14,000 D) \$22,000

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs \$28/hr; a decentralized network costs \$6/hr. How much does the startup save?

A) \$11,000 B) \$3,000 C) \$14,000 D) \$22,000

Answer: (A) – Centralized: $500 \times \$28 = \$14{,}000$. Decentralized: $500 \times \$6 = \$3{,}000$. Saving: $\$14{,}000 - \$3{,}000 = \$11{,}000$.

Q7 [Apply]. A tokenized inference service charges 0.5 tokens per query; the token price is \$2.00. What is the total cost for 1,000 queries?

A) \$500 B) \$2,000 C) \$1,000 D) \$250

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs $28/hr; a decentralized network costs $6/hr. How much does the startup save?
A) $11,000 B) $3,000 C) $14,000 D) $22,000
Answer: (A) – Centralized: $500 \times \$28 = \$14{,}000$. Decentralized: $500 \times \$6 = \$3{,}000$. Saving: $\$14{,}000 - \$3{,}000 = \$11{,}000$.

Q7 [Apply]. A tokenized inference service charges 0.5 tokens per query; the token price is $2.00. What is the total cost for 1,000 queries?
A) $500 B) $2,000 C) $1,000 D) $250
Answer: (C) – $1{,}000 \times 0.5 \times \$2.00 = \$1{,}000$ total query cost.

Q8 [Apply]. A project description reads: "Our platform uses AI to score DeFi loan applications and records each decision on Ethereum with the model hash, input hash, and zkML proof." Which dimension of the Convergence Thesis does this BEST exemplify?
A) Tokenized AI economy B) Decentralized compute
C) Verifiable AI inference with on-chain audit trail D) AI-driven oracle network

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs $28/hr; a decentralized network costs $6/hr. How much does the startup save?
A) $11,000 B) $3,000 C) $14,000 D) $22,000
Answer: (A) – Centralized: $500 \times \$28 = \$14{,}000$. Decentralized: $500 \times \$6 = \$3{,}000$. Saving: $\$14{,}000 - \$3{,}000 = \$11{,}000$.

Q7 [Apply]. A tokenized inference service charges 0.5 tokens per query; the token price is $2.00. What is the total cost for 1,000 queries?
A) $500 B) $2,000 C) $1,000 D) $250
Answer: (C) – $1{,}000 \times 0.5 \times \$2.00 = \$1{,}000$ total query cost.

Q8 [Apply]. A project description reads: "Our platform uses AI to score DeFi loan applications and records each decision on Ethereum with the model hash, input hash, and zkML proof." Which dimension of the Convergence Thesis does this BEST exemplify?
A) Tokenized AI economy B) Decentralized compute
C) Verifiable AI inference with on-chain audit trail D) AI-driven oracle network
Answer: (C) – Recording the model hash, input hash, and zkML proof on-chain makes each AI decision auditable and tamper-evident – the core of verifiable inference.

Q9 [Apply]. A content creator registers a cryptographic hash of their artwork on Ethereum. A rival platform later displays identical content without any on-chain attestation. What does this indicate?
A) The rival platform obtained a valid license B) The blockchain registration is legally unenforceable
C) The rival platform copied the content after the creator's timestamp, enabling provenance proof D) The hash collision means both parties created the content independently

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs $28/hr; a decentralized network costs $6/hr. How much does the startup save?
A) $11,000 B) $3,000 C) $14,000 D) $22,000
Answer: (A) – Centralized: $500 \times \$28 = \$14{,}000$. Decentralized: $500 \times \$6 = \$3{,}000$. Saving: $\$14{,}000 - \$3{,}000 = \$11{,}000$.

Q7 [Apply]. A tokenized inference service charges 0.5 tokens per query; the token price is $2.00. What is the total cost for 1,000 queries?
A) $500 B) $2,000 C) $1,000 D) $250
Answer: (C) – $1{,}000 \times 0.5 \times \$2.00 = \$1{,}000$ total query cost.

Q8 [Apply]. A project description reads: "Our platform uses AI to score DeFi loan applications and records each decision on Ethereum with the model hash, input hash, and zkML proof." Which dimension of the Convergence Thesis does this BEST exemplify?
A) Tokenized AI economy B) Decentralized compute
C) Verifiable AI inference with on-chain audit trail D) AI-driven oracle network
Answer: (C) – Recording the model hash, input hash, and zkML proof on-chain makes each AI decision auditable and tamper-evident – the core of verifiable inference.

Q9 [Apply]. A content creator registers a cryptographic hash of their artwork on Ethereum. A rival platform later displays identical content without any on-chain attestation. What does this indicate?
A) The rival platform obtained a valid license B) The blockchain registration is legally unenforceable
C) The rival platform copied the content after the creator's timestamp, enabling provenance proof D) The hash collision means both parties created the content independently
Answer: (C) – The on-chain timestamp precedes the rival's use; the creator can demonstrate prior existence of the content via the immutable hash record.

Q10 [Apply]. An AI oracle feeds a Bitcoin price of $50,000 to a DeFi lending protocol. The true market price is $40,000. What is the most likely consequence for the protocol?
A) Borrowers receive less collateral value than expected B) The protocol pauses automatically
C) Loans are over-collateralized, protecting lenders D) Borrowers can take out larger loans than their collateral justifies, creating bad debt

Quiz: Questions 6–10

Apply

Q6 [Apply]. A startup needs 500 GPU-hours for fine-tuning. Centralized cloud costs $28/hr; a decentralized network costs $6/hr. How much does the startup save?
A) $11,000 B) $3,000 C) $14,000 D) $22,000
Answer: (A) – Centralized: $500 \times \$28 = \$14{,}000$. Decentralized: $500 \times \$6 = \$3{,}000$. Saving: $\$14{,}000 - \$3{,}000 = \$11{,}000$.

Q7 [Apply]. A tokenized inference service charges 0.5 tokens per query; the token price is $2.00. What is the total cost for 1,000 queries?
A) $500 B) $2,000 C) $1,000 D) $250
Answer: (C) – $1{,}000 \times 0.5 \times \$2.00 = \$1{,}000$ total query cost.

Q8 [Apply]. A project description reads: "Our platform uses AI to score DeFi loan applications and records each decision on Ethereum with the model hash, input hash, and zkML proof." Which dimension of the Convergence Thesis does this BEST exemplify?
A) Tokenized AI economy B) Decentralized compute
C) Verifiable AI inference with on-chain audit trail D) AI-driven oracle network
Answer: (C) – Recording the model hash, input hash, and zkML proof on-chain makes each AI decision auditable and tamper-evident – the core of verifiable inference.

Q9 [Apply]. A content creator registers a cryptographic hash of their artwork on Ethereum. A rival platform later displays identical content without any on-chain attestation. What does this indicate?
A) The rival platform obtained a valid license B) The blockchain registration is legally unenforceable
C) The rival platform copied the content after the creator's timestamp, enabling provenance proof D) The hash collision means both parties created the content independently
Answer: (C) – The on-chain timestamp precedes the rival's use; the creator can demonstrate prior existence of the content via the immutable hash record.

Q10 [Apply]. An AI oracle feeds a Bitcoin price of $50,000 to a DeFi lending protocol. The true market price is $40,000. What is the most likely consequence for the protocol?
A) Borrowers receive less collateral value than expected B) The protocol pauses automatically
C) Loans are over-collateralized, protecting lenders D) Borrowers can take out larger loans than their collateral justifies, creating bad debt
Answer: (D) – An inflated price oracle allows borrowers to extract more value than their collateral is worth; if prices correct the protocol accumulates bad debt, as seen in the Mango Markets exploit.

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?

A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously

C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?

A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously

C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution

Answer: (B) – Concentration in a single hardware vendor creates correlated failure risk; a firmware bug or supply-chain attack can cascade across the majority of providers.

Q12 [Apply]. A DAO proposes delegating treasury management to an autonomous AI agent. Which evaluation question is MOST critical before deployment?

A) Does the agent generate higher returns than a human manager? B) Is the agent's code open source?

C) Can humans override or shut down the agent, and what are its decision boundaries? D) Does the agent use the latest language model?

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?

A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously

C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution

Answer: (B) – Concentration in a single hardware vendor creates correlated failure risk; a firmware bug or supply-chain attack can cascade across the majority of providers.

Q12 [Apply]. A DAO proposes delegating treasury management to an autonomous AI agent. Which evaluation question is MOST critical before deployment?

A) Does the agent generate higher returns than a human manager? B) Is the agent's code open source?

C) Can humans override or shut down the agent, and what are its decision boundaries? D) Does the agent use the latest language model?

Answer: (C) – Human override capability and clearly defined decision boundaries are the primary safety requirement for any autonomous agent managing real assets.

Q13 [Analyze]. Why is AI hallucination MORE dangerous when an AI model is connected to a smart contract than on a centralized chatbot platform?

A) Hallucinations occur more frequently on blockchain B) Smart contract outputs are public, so hallucinations embarrass the organization

C) Smart contract executions are irreversible and self-executing; a hallucinated output can trigger fund transfers or contract state changes that cannot be undone D) Blockchains amplify hallucinations through consensus mechanisms

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?
A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously
C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution
Answer: (B) – Concentration in a single hardware vendor creates correlated failure risk; a firmware bug or supply-chain attack can cascade across the majority of providers.

Q12 [Apply]. A DAO proposes delegating treasury management to an autonomous AI agent. Which evaluation question is MOST critical before deployment?
A) Does the agent generate higher returns than a human manager? B) Is the agent's code open source?
C) Can humans override or shut down the agent, and what are its decision boundaries? D) Does the agent use the latest language model?
Answer: (C) – Human override capability and clearly defined decision boundaries are the primary safety requirement for any autonomous agent managing real assets.

Q13 [Analyze]. Why is AI hallucination MORE dangerous when an AI model is connected to a smart contract than on a centralized chatbot platform?
A) Hallucinations occur more frequently on blockchain B) Smart contract outputs are public, so hallucinations embarrass the organization
C) Smart contract executions are irreversible and self-executing; a hallucinated output can trigger fund transfers or contract state changes that cannot be undone D) Blockchains amplify hallucinations through consensus mechanisms
Answer: (C) – Centralized platforms can roll back incorrect outputs; blockchain transactions are immutable, so a hallucinated decision executed on-chain may permanently alter state or move funds.

Q14 [Analyze]. Compare centralized AI (e.g., OpenAI API) and decentralized AI (e.g., Bittensor). What is the KEY structural difference regarding censorship resistance?
A) Decentralized AI is always more accurate B) Centralized AI is cheaper to run
C) Centralized AI can restrict or modify access based on policy; decentralized AI permissioned by token economics has no single point of control to censor D) There is no meaningful difference in censorship resistance

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?

A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously

C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution

Answer: (B) – Concentration in a single hardware vendor creates correlated failure risk; a firmware bug or supply-chain attack can cascade across the majority of providers.

Q12 [Apply]. A DAO proposes delegating treasury management to an autonomous AI agent. Which evaluation question is MOST critical before deployment?

A) Does the agent generate higher returns than a human manager? B) Is the agent's code open source?

C) Can humans override or shut down the agent, and what are its decision boundaries? D) Does the agent use the latest language model?

Answer: (C) – Human override capability and clearly defined decision boundaries are the primary safety requirement for any autonomous agent managing real assets.

Q13 [Analyze]. Why is AI hallucination MORE dangerous when an AI model is connected to a smart contract than on a centralized chatbot platform?

A) Hallucinations occur more frequently on blockchain B) Smart contract outputs are public, so hallucinations embarrass the organization

C) Smart contract executions are irreversible and self-executing; a hallucinated output can trigger fund transfers or contract state changes that cannot be undone D) Blockchains amplify hallucinations through consensus mechanisms

Answer: (C) – Centralized platforms can roll back incorrect outputs; blockchain transactions are immutable, so a hallucinated decision executed on-chain may permanently alter state or move funds.

Q14 [Analyze]. Compare centralized AI (e.g., OpenAI API) and decentralized AI (e.g., Bittensor). What is the KEY structural difference regarding censorship resistance?

A) Decentralized AI is always more accurate B) Centralized AI is cheaper to run

C) Centralized AI can restrict or modify access based on policy; decentralized AI permissioned by token economics has no single point of control to censor D) There is no meaningful difference in censorship resistance

Answer: (C) – A single company controls access to a centralized API and can block users or alter outputs; decentralized networks distribute control so no entity can unilaterally censor.

Q15 [Analyze]. Fetch.ai, SingularityNET, and Ocean Protocol merged into the ASI Alliance. What strategic problem does this merger address?

A) All three companies were unprofitable and needed cost savings B) Regulators required consolidation

C) Individually, each protocol lacked scale to compete with hyperscaler AI platforms; the merger creates a larger token market cap and combined network effect to attract developers and compute D) The merger was required to list on major exchanges

Quiz: Questions 11–15

Apply / Analyze

Q11 [Apply]. A decentralized compute network has 100 providers; 80 run NVIDIA A100s and 20 run older NVIDIA V100s. One provider uses custom TPUs. Which risk does this illustrate?
A) Sybil attack risk B) Hardware monoculture – a driver vulnerability in NVIDIA GPUs could disable 80% of capacity simultaneously
C) Regulatory risk from using proprietary chips D) Latency risk from geographic distribution
Answer: (B) – Concentration in a single hardware vendor creates correlated failure risk; a firmware bug or supply-chain attack can cascade across the majority of providers.

Q12 [Apply]. A DAO proposes delegating treasury management to an autonomous AI agent. Which evaluation question is MOST critical before deployment?
A) Does the agent generate higher returns than a human manager? B) Is the agent's code open source?
C) Can humans override or shut down the agent, and what are its decision boundaries? D) Does the agent use the latest language model?
Answer: (C) – Human override capability and clearly defined decision boundaries are the primary safety requirement for any autonomous agent managing real assets.

Q13 [Analyze]. Why is AI hallucination MORE dangerous when an AI model is connected to a smart contract than on a centralized chatbot platform?
A) Hallucinations occur more frequently on blockchain B) Smart contract outputs are public, so hallucinations embarrass the organization
C) Smart contract executions are irreversible and self-executing; a hallucinated output can trigger fund transfers or contract state changes that cannot be undone D) Blockchains amplify hallucinations through consensus mechanisms
Answer: (C) – Centralized platforms can roll back incorrect outputs; blockchain transactions are immutable, so a hallucinated decision executed on-chain may permanently alter state or move funds.

Q14 [Analyze]. Compare centralized AI (e.g., OpenAI API) and decentralized AI (e.g., Bittensor). What is the KEY structural difference regarding censorship resistance?
A) Decentralized AI is always more accurate B) Centralized AI is cheaper to run
C) Centralized AI can restrict or modify access based on policy; decentralized AI permissioned by token economics has no single point of control to censor D) There is no meaningful difference in censorship resistance
Answer: (C) – A single company controls access to a centralized API and can block users or alter outputs; decentralized networks distribute control so no entity can unilaterally censor.

Q15 [Analyze]. Fetch.ai, SingularityNET, and Ocean Protocol merged into the ASI Alliance. What strategic problem does this merger address?
A) All three companies were unprofitable and needed cost savings B) Regulators required consolidation
C) Individually, each protocol lacked scale to compete with hyperscaler AI platforms; the merger creates a larger token market cap and combined network effect to attract developers and compute D) The merger was required to list on major exchanges
Answer: (C) – Fragmented decentralized AI projects struggle to match the network effects of OpenAI or Google; pooling tokens, users, and developer ecosystems creates a more credible alternative.

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial
Answer: (C) – When one supplier dominates, geopolitical restrictions (e.g., US export controls on A100s to China) or supply shocks propagate across the entire decentralized ecosystem simultaneously.

Q17 [Analyze]. A zkML proof is generated for a credit-scoring model that returned "reject" for a loan applicant. Why is this useful to the applicant even without revealing the model?
A) The proof allows the applicant to reverse the decision automatically B) The proof reveals the model weights so the applicant can audit bias
C) The proof confirms the decision was produced by the claimed model running correctly on the claimed inputs, enabling third-party auditors to verify fairness without the bank disclosing proprietary IP D) The proof encrypts the applicant's data so the bank cannot see it

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial
Answer: (C) – When one supplier dominates, geopolitical restrictions (e.g., US export controls on A100s to China) or supply shocks propagate across the entire decentralized ecosystem simultaneously.

Q17 [Analyze]. A zkML proof is generated for a credit-scoring model that returned "reject" for a loan applicant. Why is this useful to the applicant even without revealing the model?
A) The proof allows the applicant to reverse the decision automatically B) The proof reveals the model weights so the applicant can audit bias
C) The proof confirms the decision was produced by the claimed model running correctly on the claimed inputs, enabling third-party auditors to verify fairness without the bank disclosing proprietary IP D) The proof encrypts the applicant's data so the bank cannot see it
Answer: (C) – zkML separates execution integrity from secrecy: the applicant (or regulator) can verify *which* model ran and *that* it ran correctly, without the bank exposing its scoring algorithm.

Q18 [Analyze]. Blockchain content provenance works well for newly created content. Why does it NOT solve the problem of existing deepfakes already circulating online?
A) Blockchains are too slow to process existing content B) Deepfakes are registered on alternative blockchains
C) Provenance systems establish origin going forward; they cannot retroactively authenticate content created before registration, and deepfakes already spread virally before provenance tools are adopted D) Deepfakes cannot be detected by hash comparison

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial
Answer: (C) – When one supplier dominates, geopolitical restrictions (e.g., US export controls on A100s to China) or supply shocks propagate across the entire decentralized ecosystem simultaneously.

Q17 [Analyze]. A zkML proof is generated for a credit-scoring model that returned "reject" for a loan applicant. Why is this useful to the applicant even without revealing the model?
A) The proof allows the applicant to reverse the decision automatically B) The proof reveals the model weights so the applicant can audit bias
C) The proof confirms the decision was produced by the claimed model running correctly on the claimed inputs, enabling third-party auditors to verify fairness without the bank disclosing proprietary IP D) The proof encrypts the applicant's data so the bank cannot see it
Answer: (C) – zkML separates execution integrity from secrecy: the applicant (or regulator) can verify *which* model ran and *that* it ran correctly, without the bank exposing its scoring algorithm.

Q18 [Analyze]. Blockchain content provenance works well for newly created content. Why does it NOT solve the problem of existing deepfakes already circulating online?
A) Blockchains are too slow to process existing content B) Deepfakes are registered on alternative blockchains
C) Provenance systems establish origin going forward; they cannot retroactively authenticate content created before registration, and deepfakes already spread virally before provenance tools are adopted D) Deepfakes cannot be detected by hash comparison
Answer: (C) – Provenance is prospective: registering future content helps, but the vast corpus of already-circulating manipulated media has no prior on-chain record to invalidate.

Q19 [Evaluate]. A project claims "AI-powered blockchain analytics with our governance token." Apply the 5-question evaluation framework. What is the MOST appropriate assessment?
A) Strong project – combines two hot sectors B) Promising – governance tokens always add value
C) Requires more information – the claim does not specify what AI does, why blockchain is needed, what the token incentivises, or whether AI adds unique value over a centralised database
D) Weak project – AI and blockchain are incompatible

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial
Answer: (C) – When one supplier dominates, geopolitical restrictions (e.g., US export controls on A100s to China) or supply shocks propagate across the entire decentralized ecosystem simultaneously.

Q17 [Analyze]. A zkML proof is generated for a credit-scoring model that returned "reject" for a loan applicant. Why is this useful to the applicant even without revealing the model?
A) The proof allows the applicant to reverse the decision automatically B) The proof reveals the model weights so the applicant can audit bias
C) The proof confirms the decision was produced by the claimed model running correctly on the claimed inputs, enabling third-party auditors to verify fairness without the bank disclosing proprietary IP D) The proof encrypts the applicant's data so the bank cannot see it
Answer: (C) – zkML separates execution integrity from secrecy: the applicant (or regulator) can verify *which* model ran and *that* it ran correctly, without the bank exposing its scoring algorithm.

Q18 [Analyze]. Blockchain content provenance works well for newly created content. Why does it NOT solve the problem of existing deepfakes already circulating online?
A) Blockchains are too slow to process existing content B) Deepfakes are registered on alternative blockchains
C) Provenance systems establish origin going forward; they cannot retroactively authenticate content created before registration, and deepfakes already spread virally before provenance tools are adopted D) Deepfakes cannot be detected by hash comparison
Answer: (C) – Provenance is prospective: registering future content helps, but the vast corpus of already-circulating manipulated media has no prior on-chain record to invalidate.

Q19 [Evaluate]. A project claims "AI-powered blockchain analytics with our governance token." Apply the 5-question evaluation framework. What is the MOST appropriate assessment?
A) Strong project – combines two hot sectors B) Promising – governance tokens always add value
C) Requires more information – the claim does not specify what AI does, why blockchain is needed, what the token incentivises, or whether AI adds unique value over a centralised database
D) Weak project – AI and blockchain are incompatible
Answer: (C) – The 5 questions are: Does AI add value here? Does blockchain add value? What does the token incentivise? Is decentralisation necessary? What is the real cryptoeconomic model? None are answered by the claim.

Q20 [Evaluate]. A government proposes requiring all AI models used in financial services to publish zkML proofs on a public blockchain. Evaluate this proposal on privacy, cost, and feasibility.
A) Fully sound – zkML proofs protect all private data while ensuring transparency B) Partially sound – proofs protect model IP but on-chain transaction data may still expose client patterns; proof generation is computationally expensive at scale; standardising proof systems across model architectures remains unsolved
C) Unsound – zkML is incompatible with financial regulation D) Fully unsound – zero-knowledge proofs cannot be generated for neural networks

Quiz: Questions 16–20

Analyze / Evaluate

Q16 [Analyze]. GPU manufacturing is dominated by NVIDIA (>80% data-centre share). How does this concentration affect decentralized compute networks?
A) It has no effect because decentralized networks run on consumer CPUs B) It is beneficial because standardisation reduces integration cost
C) It creates a correlated supply risk: shortages, price shocks, or export controls on NVIDIA chips simultaneously constrain all providers D) NVIDIA actively supports decentralised networks so concentration is beneficial
Answer: (C) – When one supplier dominates, geopolitical restrictions (e.g., US export controls on A100s to China) or supply shocks propagate across the entire decentralized ecosystem simultaneously.

Q17 [Analyze]. A zkML proof is generated for a credit-scoring model that returned "reject" for a loan applicant. Why is this useful to the applicant even without revealing the model?
A) The proof allows the applicant to reverse the decision automatically B) The proof reveals the model weights so the applicant can audit bias
C) The proof confirms the decision was produced by the claimed model running correctly on the claimed inputs, enabling third-party auditors to verify fairness without the bank disclosing proprietary IP D) The proof encrypts the applicant's data so the bank cannot see it
Answer: (C) – zkML separates execution integrity from secrecy: the applicant (or regulator) can verify *which* model ran and *that* it ran correctly, without the bank exposing its scoring algorithm.

Q18 [Analyze]. Blockchain content provenance works well for newly created content. Why does it NOT solve the problem of existing deepfakes already circulating online?
A) Blockchains are too slow to process existing content B) Deepfakes are registered on alternative blockchains
C) Provenance systems establish origin going forward; they cannot retroactively authenticate content created before registration, and deepfakes already spread virally before provenance tools are adopted D) Deepfakes cannot be detected by hash comparison
Answer: (C) – Provenance is prospective: registering future content helps, but the vast corpus of already-circulating manipulated media has no prior on-chain record to invalidate.

Q19 [Evaluate]. A project claims "AI-powered blockchain analytics with our governance token." Apply the 5-question evaluation framework. What is the MOST appropriate assessment?
A) Strong project – combines two hot sectors B) Promising – governance tokens always add value
C) Requires more information – the claim does not specify what AI does, why blockchain is needed, what the token incentivises, or whether AI adds unique value over a centralised database
D) Weak project – AI and blockchain are incompatible
Answer: (C) – The 5 questions are: Does AI add value here? Does blockchain add value? What does the token incentivise? Is decentralisation necessary? What is the real cryptoeconomic model? None are answered by the claim.

Q20 [Evaluate]. A government proposes requiring all AI models used in financial services to publish zkML proofs on a public blockchain. Evaluate this proposal on privacy, cost, and feasibility.
A) Fully sound – zkML proofs protect all private data while ensuring transparency B) Partially sound – proofs protect model IP but on-chain transaction data may still expose client patterns; proof generation is computationally expensive at scale; standardising proof systems across model architectures remains unsolved
C) Unsound – zkML is incompatible with financial regulation D) Fully unsound – zero-knowledge proofs cannot be generated for neural networks
Answer: (B) – zkML offers meaningful transparency benefits but faces three constraints: metadata leakage from on-chain records, high proof-generation cost for large models, and lack of standardised tooling across neural architectures – making broad mandates premature.