

Bicsi[®]
ENDORSED EVENT

ICT SUMMIT COLOMBIA 2024

Explorando Inteligencia, Digitalización y Sustentabilidad
para el Futuro del ICT.

M A Y O

MARTES 28 Y MIÉRCOLES 29

Agora Bogotá Centro de Convenciones
Bogotá | Colombia

ORGANIZA:

LATAM  **RED**

How new BICSI & TIA data Centers Support AI & other new technologies

Jonathan Jew

J&M Consultants, Inc

jew@j-and-m.com

Editor ANSI/TIA-942-C Data Center Standard

Chair BICSI Data Center Design Working Group

A New Revolution

- Artificial Intelligence will be as great an impact on people and society as computers





新春大吉


中美銀行
S. M. G. INC.

大華行
G. INC.

Pay here
Pay by License Plate
72009003

Artificial Intelligence (AI) uses GPUs

- Data centers with graphics processing unit (GPUs) servers perform the learning used for AI
- GPU processors include Nvidia, AMD, Intel and others



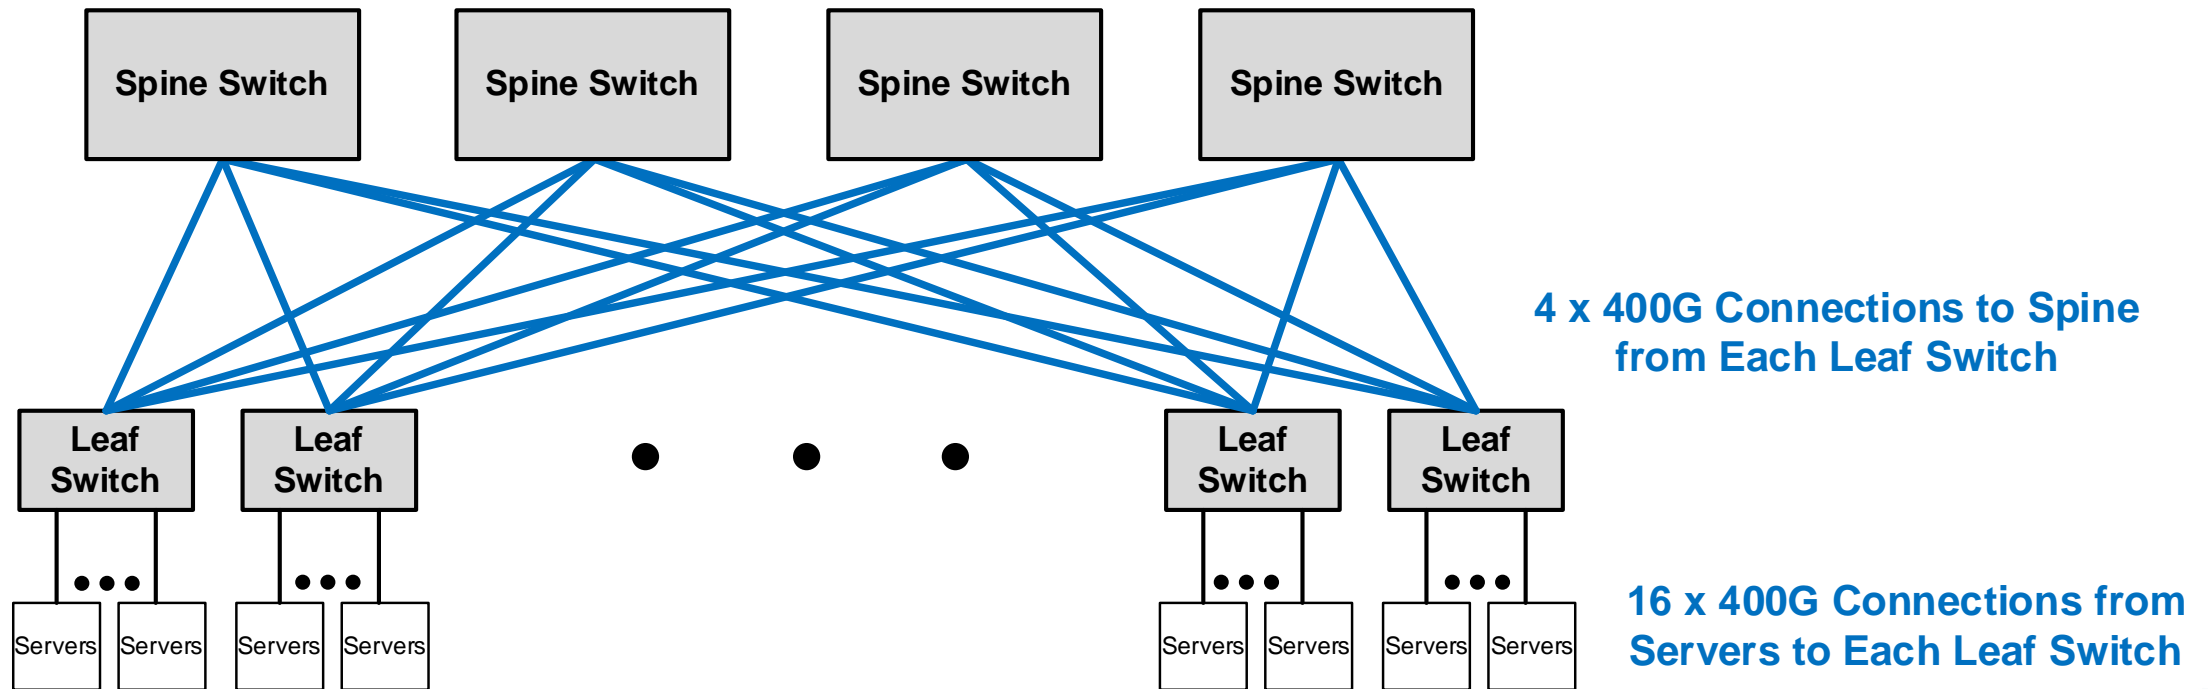
Common Configuration of GPU Servers

- 6 to 8 ports 400G Ethernet or InfiniBand for GPU network (non-blocking)
- 2 ports 400G Ethernet or InfiniBand for storage network may be non-blocking or over-subscribed
- 2 ports 100G for in-band management typically over-subscribed
- 1 port 1 Gbps Ethernet out-of-band (OOB) management (e.g., IPMI, iDRAC, ILO) over-subscribed



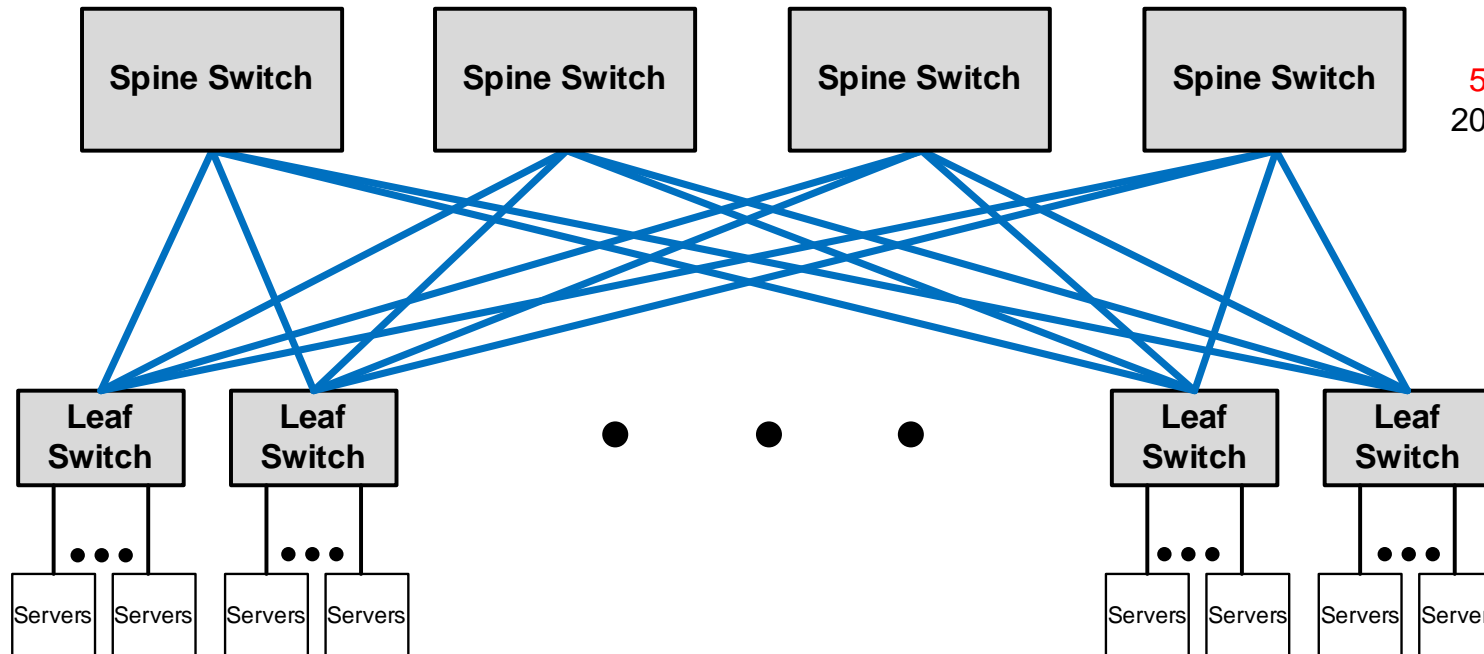
4:1 Over-Subscribed Network Example

Storage & In-Band Management Networks may be Over-Subscribed and use 100G – 400G ports
(not all ports transmit at full speed concurrently)



Non-Blocking Network Example

GPU Networks are Non-Blocking Commonly 400G with New Deployments
(all ports may transmit at 400G at the same time)



512 server ports (64 servers) max with 32-port switches
2048 server ports (256 servers) max with 64-port switches

16 x 400G Connections to Spine
from Each Leaf Switch
All Leafs Connected to All Spines

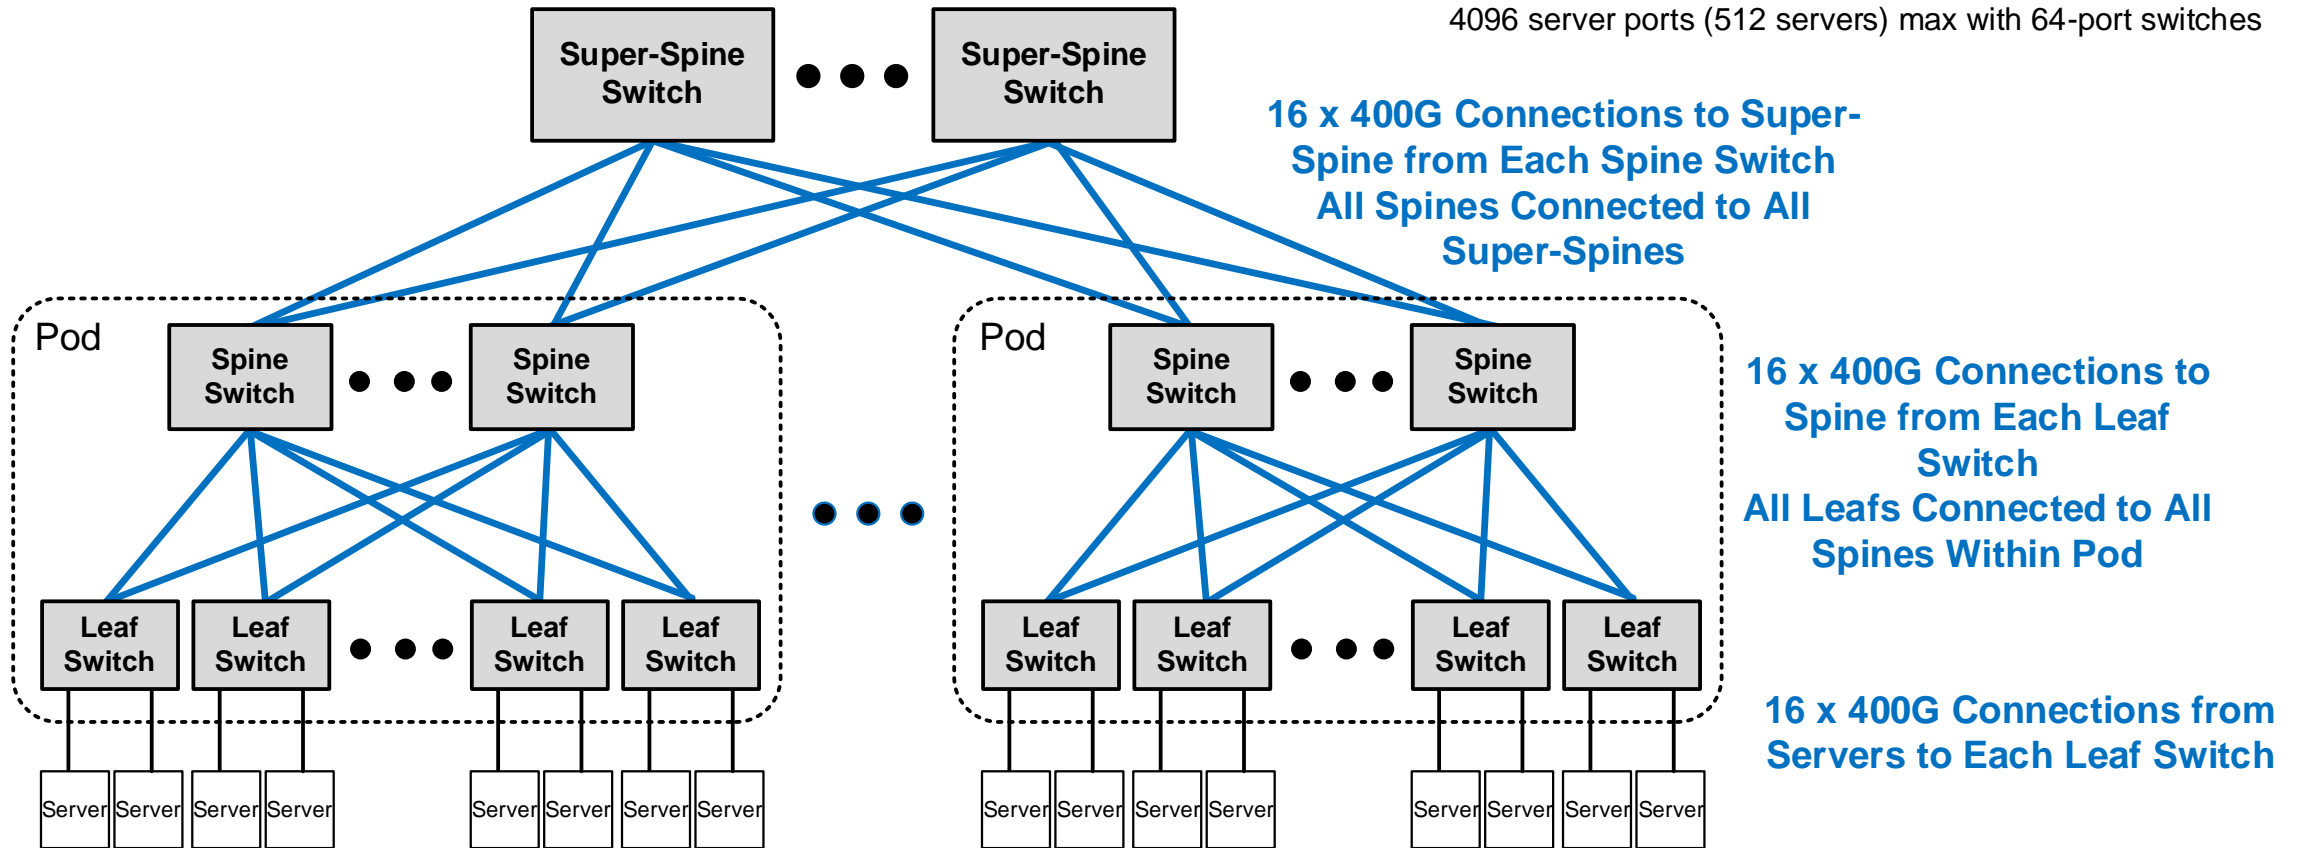
16 x 400G Connections from
Servers to Each Leaf Switch

Leaf Switches have $\frac{1}{2}$ of ports for servers and $\frac{1}{2}$ of ports to spine

Super-Spine to Increase Network Size

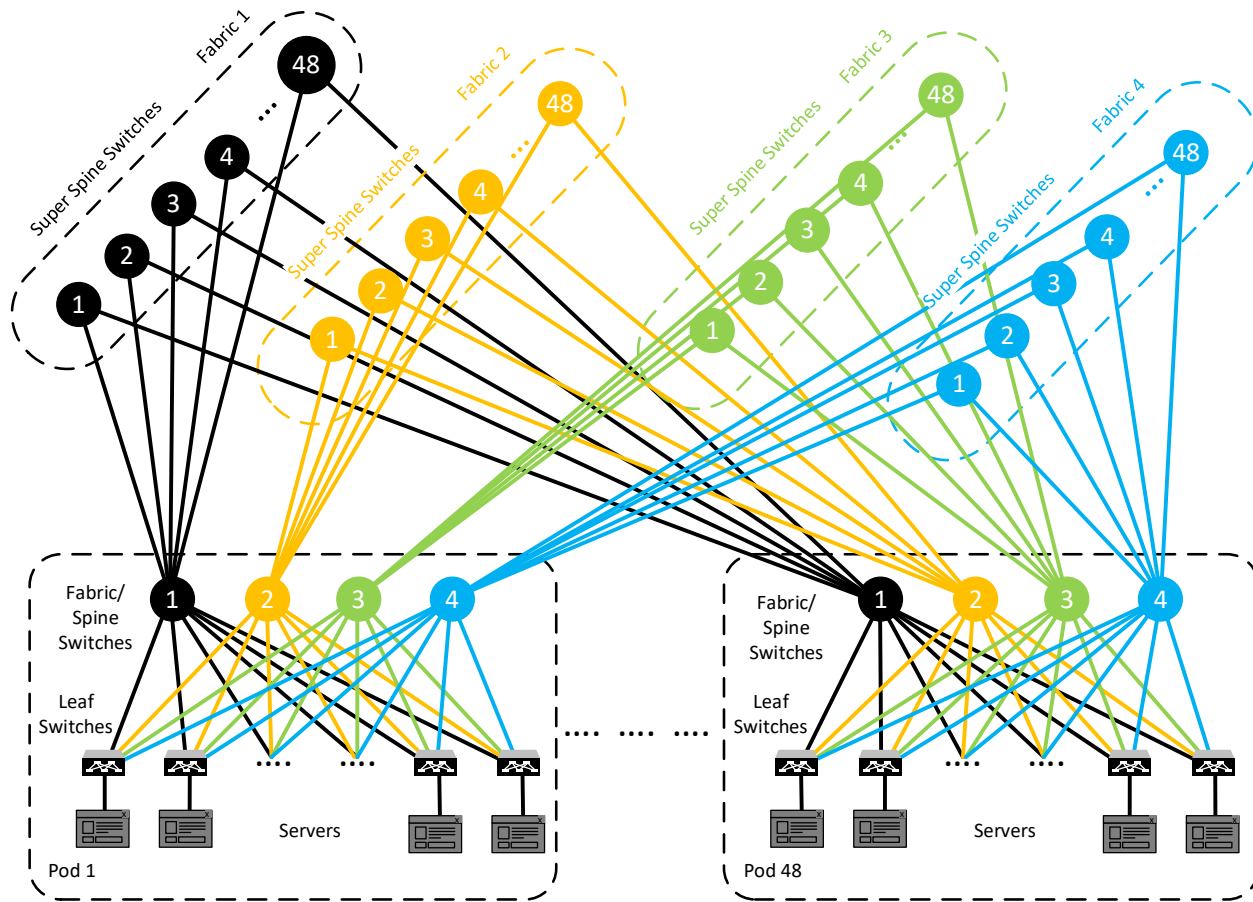
1024 server ports (128 servers) max with 32-port switches

4096 server ports (512 servers) max with 64-port switches



Leaf & Spine Switches have ½ of ports going down 1-level and ½ of ports going up one level

Hyperscale Fabric for Larger AI Networks




- Leaf switches connected to a spine/fabric switch for each of the fabrics
- Each spine/fabric switch connected to all super-spine switches in its fabric
- Leaf & spine switches have $\frac{1}{2}$ of ports up 1-level and $\frac{1}{2}$ of ports down 1 level
- The network can grow larger by deploying higher port-count spine & super-spine switches or by adding fabrics
- 4 fabric network with 64-port spines & super-spines has 8,192 server ports (1024 servers)
- 4 fabric network with 128-port spines & super-spines has 32,768 server ports (4096 servers)

AI Has High Networking & Cabling Demands

- **Very high speeds - 400G & 800G** ports currently
- Non-blocking networks have **lots of connections**
- For 1,024 servers (with 8 x 400G ports each – 8,192 total)
 - 16,384 400G leaf switch ports (**typically direct cables**)
 - 16,384 400G spine switch ports (**should use structured cabling to spines in intermediate or main distributor**)
 - 8,192 400G super-spine switch ports (**should use structured cabling to super-spines in main distributor**)
- Commonly 4 additional ports per server, 2 each for storage & in-band networks at 400G to 100G may be non-blocking or over subscribed
- 400G cables typically use 8-fibers with angled-MPOs

2024 Announcements regarding 2025 Products

- 800G NICs on GPU servers
- 144-port 4U 800G InfiniBand switch & 64-port 2U 800G Ethernet switch
- Transceivers are 800G with either
 - Two 400GBASE-DR4 ports (SMF MPO12/APC with 8 fibers, 500m)
 - Two 400GBASE-FR4 ports (SMF LC/UPC with 2 fibers, 2km)
- Current 800G switches are 32-port 1U switches



- Mar 2026 IEEE 802.3dj
- 800G over 8 SMF
 - 1.6 TB over 16 SMF (e.g., two MPOs each at 800G)



AI Data Center Concerns

- Highly connected at very high speeds
- Ultra-high-density connectivity to conserve cabinet space
- May use single-mode optical fiber for migration to 800G & 1.6T and distances greater than 100 m
- High power consumption (30 kw+ per cabinet air cooled & higher with liquid cooling)
- Deploy quickly to support new AI needs


AI Applications tied to IOT & Edge

- IoT for sensors used by AI to gather information about the physical world
- Edge Data Centers to
 - consolidate & summarize IoT data,
 - perform local processing (e.g., for an automated factory),
 - reduce delay



Complementary Standards

- ANSI/BICSI-002 Data Center Design & Implementation Best Practices
 - Covers a wide range subjects including electrical, mechanical, architectural, security, and telecom
 - Complements the local data center telecom standard
- ANSI/TIA-942 Telecommunications Infrastructure standard for Data Centers
 - The data center telecom standard for the US
 - Widely used globally



Updates to ANSI/BICSI 002-2024

- Added content on immersion cooling and liquid cooling systems
- Cabinet cooling systems redundancy should be on UPS and match the desired Class (tier/rating)
- Both of these are important considerations to support higher power and cooling requirements of AI servers

Updates to ANSI/BICSI 002-2024 (continued)

- Added single-pair and added broadband coaxial cable as allowed types of cables
 - Single-pair should be widely used for IoT, broadband coax is used for some applications
- Removed recommendation to use LC and MPO in distributors
 - More flexibility for selection of Very Small Form Factor optical fiber connectors in high density distributors. Very high density is a need for cabling for AI
- Minimum width of cabinets in distributors is 800 mm
 - 800 mm wide cabinets are needed in distributors to properly manage cables

Updates to ANSI/TIA-942-C

- Added single balanced twisted-pair cable as a recognized type of cable for horizontal cabling
- Fiber connectors may be any TIA-568.3 compliant connector outside the EO in the server cabinet
- Added requirement of a minimum of 800 mm wide cabinets in distributors
- Similar changes made in BICSI-002 as some people worked on both standards

Updates to ANSI/TIA-942-C (continued)

- Incorporated Edge Data Center Addendum with updates to requirements, recommendations, and rating system for micro edge data centers
- Added recommendations for managing micro edge data centers provided by ASHRAE
 - micro edge data center: A small data center in a pre-manufactured enclosure that is capable of being remotely monitored and located at the network edge supporting equipment and applications.
 - Different from an edge data center at the edge of a network, but that may be in a normal data center.

Updates to ANSI/TIA-942-C (continued)

- Updated temperature and humidity guidelines to require the recommended temperature and humidity envelope per ASHRAE TC 9.9
 - Change made at the recommendation of the founders of ASHRAE TC 9.9. Working outside the recommended ranges requires careful study.
- Added informative annex on Immersion Cooling Considerations
 - Cooling is important in AI data centers and immersion cooling is one of the new approaches.
- Multiple updates to Annexes to simplify use of ANSI/TIA-942 for rating data centers
 - New AI data centers are commonly being built in 3rd party data centers. Rating and certification are good tools to ensure that they will meet availability goals.

AI is another powerful tool that will change society as computers did

- The AI revolution is only starting
- AI revolution presents enormous opportunities for people like us in the ICT industry that design and install and manage infrastructure used by AI
- It also presents great opportunities for organizations that provide products (e.g., cabinets, networking hardware, servers, and cabling) used by AI
- ANSI/BICSI-002-2024 and ANSI/TIA-942-C are great standards to help everyone in ICT involved in data centers (AI or not)

Questions?

- Jonathan Jew - jew@j-and-m.com
- Editor of ANSI/TIA-942-C
- Chair of working group for ANSI/BICSI 002-2024