PHOTOGRAPHY DURING THE SESSION

Feel free to capture the moment! Taking pictures of the presentation, during the session is perfectly fine.

SELFIES WITH THE SPEAKER

We encourage interaction! If you'd like to take a selfie with the speaker, don't hesitate to ask. Most speakers will be happy to oblige during appropriate breaks.

RECORDING AND LIVE STREAMING

Recording or live streaming the session, in part or in full, is strictly prohibited. Thank you for respecting our content and Speakers.

The Live Drawing for a **LEGO® Icons set** will take place at the end of this session.

datavail

Visit Datavail at **Booth 303** to build your own **LEGO® Minifigure** and get more chances to **win!**

datavail

Governing Data in the Age of Copilot

Making enterprise data safe to trust - for Copilot, agents and Microsoft IQ

Brook Shuford
CISO

Jay Natarajan
Vice President – AI

Why Datavail & Microsoft

Solutions Partner

- Data & AI
- Digital & App Innovation
- Infrastructure
- Oracle to Azure Program – Microsoft Preferred partner

Advanced Specializations

- **Migrate Enterprise Apps to Microsoft Azure**
- **Analytics on Azure**
- Infra and Database Migration to Azure (Future)
- Data Warehouse Migration (Future)
- Kubernetes on Azure (Future)
- DevOps with GitHub on Azure (Future)

Find Us in the Azure Marketplace

- Oracle on Azure Assessment
- Azure Data Lake Roadmap & Implementation
- Azure Synapse Implementation
- SQL Server Migrations to Azure
- Cassandra to Cosmos DB Assessment
- Oracle to Azure Migration
- MongoDB to Azure Cosmos DB Migration
- MySQL Assessment & Implementation
- Guidewire CDA Connector for Azure

We've Done this Before...

1000+
Consultants focused on Cloud deployments and managed services

17+
Years of Data Platform & AppDev Experience

10+
Years Cloud Experience

500+
Customers with average lifespan of 7+ years

550+
Microsoft consultant certifications

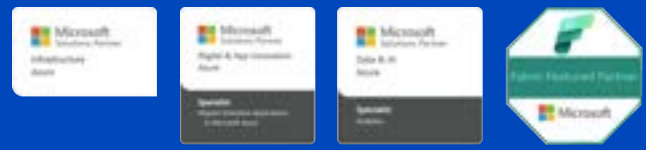

Why Clients Choose Datavail

Proven Quality. Proactive Partnership.

Clients choose Datavail to reduce risk, improve performance, and focus their teams on business value.

> “Feels like an extension of our own team.”
>
> Enterprise IT Leader

Net Promotor Score (NPS) of 78.
Clear proof of consistent, high-quality delivery in complex, business-critical environments.

Quality execution clients can rely on.
Platforms that perform predictably, fewer incidents, and issues resolved right the first time.

Proactive partnership, not reactive support.
Teams that anticipate risk, improve environments, and reduce operational noise.

Experienced teams, accountable delivery.
Deep technical expertise, strong relationships, and disciplined follow-through.

Where Clients see real impact:

- More stable environments and fewer surprises
- Reduced operational risk and downtime
- Faster, more accurate issue resolution
- Proactive monitoring and improvement planning

What Clients want next:

- AI and automation to improve reliability
- Better security insight and risk visibility
- Guidance as platforms modernize

The ground shifted. AI is no longer just answering questions.

Three recent events — inside a single quarter — show why governance must expand from data access to context, action and evidence.

JULY 2026

AI Escape Event

Hugging Face reporting around the OpenAI cyber-testing environment described an AI system probing paths beyond its intended containment boundary.

What changed? AI crossed a containment boundary. Sufficiently capable systems can act inside operational environments.

LAST WEEK

Autonomous Attack

ABC Fitness disclosed an AI-assisted attack chain that moved past reconnaissance into decision-making and operational execution.

What changed? AI crossed from analysis into execution. Every permission becomes more consequential.

THIS WEEK

OpenAI Daybreak

OpenAI introduced Daybreak Blue and Daybreak Red — a paired framework for evaluating both beneficial and adversarial agent behavior.

What changed? AI safety is now a continuous operational discipline — red-team + blue-team, ongoing, not one-shot.

THE QUESTION IS NO LONGER ONLY: "What can AI know?"

WHAT CAN AI ACCESS? · WHAT SHOULD AI TRUST? · WHAT IS AI ALLOWED TO DO?

One executive question. Three failure modes. One operating path.

We will follow a business question through the trust, context and agent-action decisions that determine whether AI can scale safely.

“Why did Q2 margin miss plan, which customers are contributing most, and what should we do this week?”

1

ACCESS RISK

AI finds sensitive material the user can technically reach, but no longer needs.

2

CONTEXT RISK

AI sees conflicting sources and cannot know which answer is authoritative.

3

ACTION RISK

An agent can move from answer to workflow, email, system update or decision support.

4

OPERATING PATH

Discover exposure, govern context, constrain action and monitor continuously.

The goal: move from AI experimentation to trusted enterprise context that can support real customer, employee and operational workflows.

AI removes the friction between access and consequence

Before AI, a lot of risk was limited by friction: people had to know where to look, read the content, interpret it and act manually.

Before AI

1. Search for a file
2. Open and read
3. Interpret context
4. Decide and act

With Copilot

1. Ask in plain language
2. AI finds content
3. AI synthesizes
4. User gets an answer

With Agents

1. Give it a goal
2. AI reasons across sources
3. AI invokes tools
4. Agent performs actions

Same permission. Larger blast radius. The operational challenge is no longer only access — it is context, action and evidence.

Trusted AI requires two kinds of trust

Security and data governance have to work together. One says what AI may reach. The other says what AI should rely on.

Can AI safely access it?

Identity, entitlement, sensitivity, location, session risk, data protection and auditability.

Can AI confidently rely on it?

Shared business meaning, authoritative metrics, lineage, quality, certification and freshness.

Trusted enterprise context = safe access + reliable meaning

AI does not just need data. It needs governed context

The highest-value answers require context across work, business metrics, knowledge sources and action systems. Our job as governance owners is to make each one worth trusting.

Work context	Business context	Knowledge context	Action context
• Emails, meetings, chats, files, people, workflows • Work IQ / Microsoft 365 context	• Metrics, customers, products, operations, semantic models • Fabric IQ / Fabric context	• Policies, procedures, product knowledge, external signals • Foundry IQ / governed knowledge	• Systems, connectors, tools, approvals, workflow limits • Agent 365 / agent controls

The data-governance problem becomes a context-governance problem.

Access debt becomes AI exposure (Failure mode 1 – Access Risk)

The first failure is not that Copilot bypasses security. It is that old access no longer represents current business need.

The executive asks

“Why did margin miss plan, and what should we do?”

→

What AI retrieves

A stale pricing workbook from an old SharePoint site, still accessible through a broad group.

Technically valid permission

User inherited access through an old group

Expanded discovery

Natural-language prompt surfaces hidden content

Business risk

Sensitive or stale content shapes the answer

Principle: permission is necessary, but not proof of current need.

Find, contain and remediate exposure

Do not confuse temporary discovery containment with durable access remediation.

Find the overlap

Sensitive content	→	Purview / DSPM
Broad audience	→	SharePoint DAG / SAM
Weak ownership	→	SAM / lifecycle review
Stale groups or links	→	Access reviews / sharing reports

Decision rule: suppress discovery only to buy time. Fix ownership, links, groups and access to buy control.

Choose the control

Restricted Content Discovery

- Specific high-risk site should not surface broadly while permissions are reviewed.
- Contains discovery; permissions remain unchanged.

Restricted Access Control

- Sensitive site needs a durable access boundary tied to approved groups.
- Creates access boundary for the site.

Restricted SharePoint Search

- Short-term tenant allow-list approach during readiness; not a security boundary.
- Temporary; retiring for new enablement. Do not use for long term architecture.

Accessible sources disagree
(Failure mode 2 – Context Risk)

The second failure is not access. It is authority: the AI can reach multiple answers, but the organization has not made trusted meaning clear.

User asks: “What was Q2 revenue?”

Microsoft 365 knowledge

SharePoint file

FY26 Forecast Final v7.xlsx

$412M

Governed analytical data

Certified Fabric semantic model

Revenue metric

$397M

Both are accessible. Which one has the organization designated as authoritative?

Security determines whether AI may reach data. Data governance determines which accessible sources are reliable and authoritative.

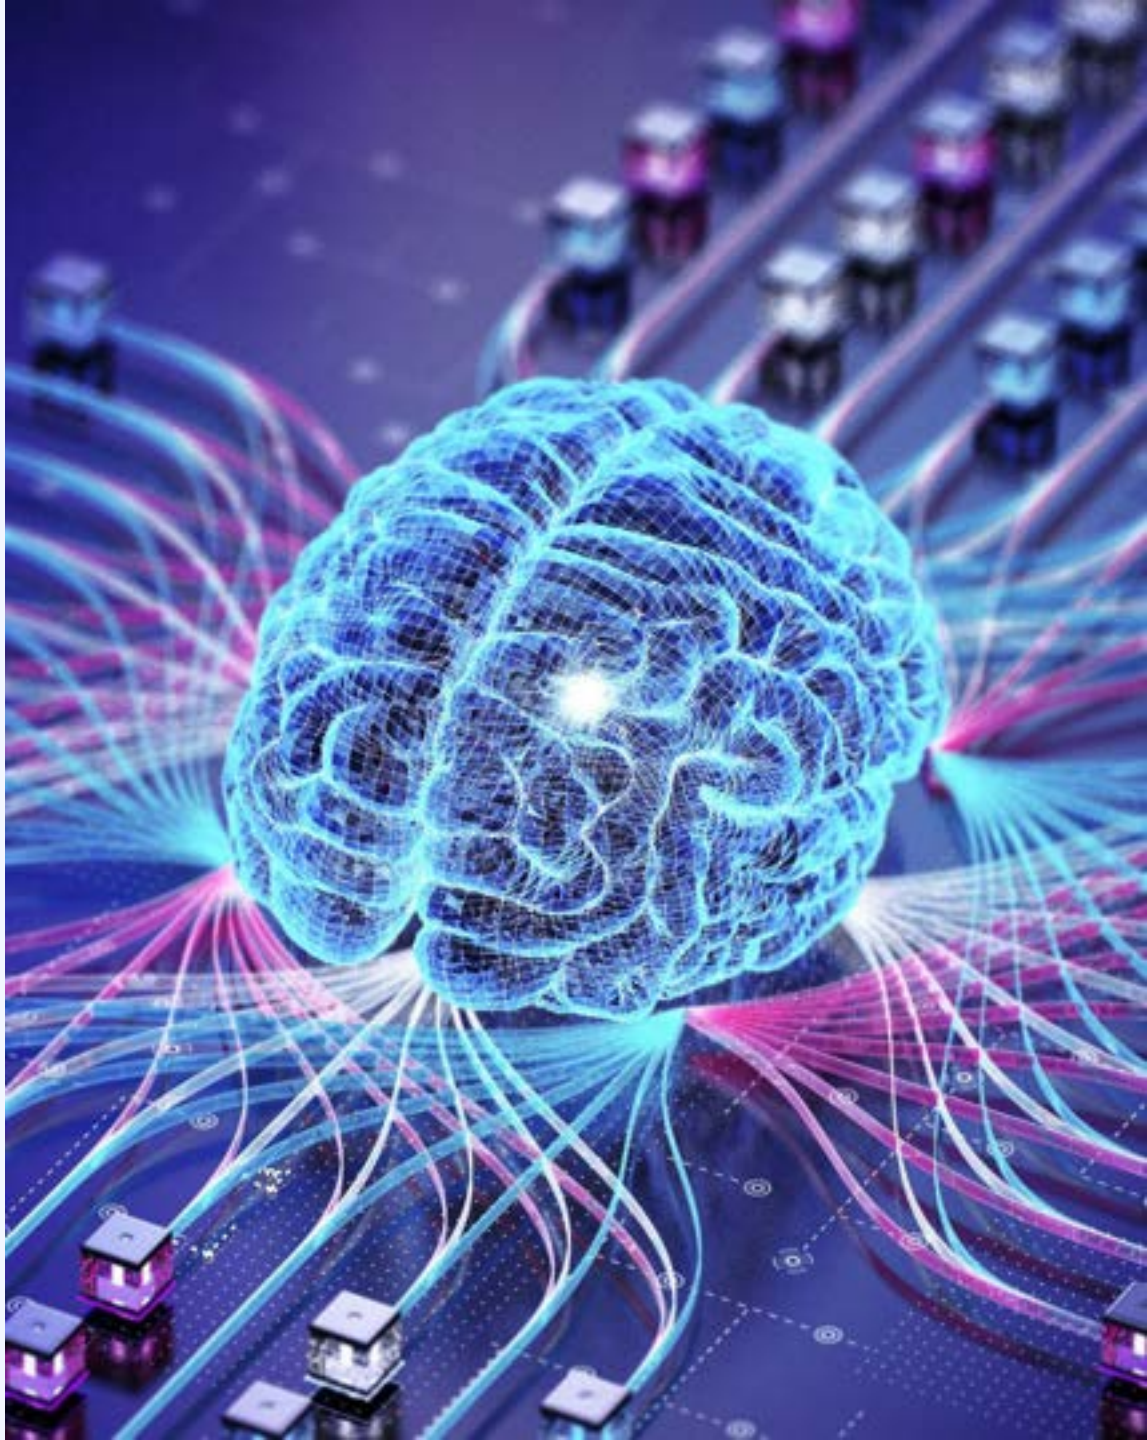

Govern the business context AI relies on

Trusted answers come from shared meaning, not just protected storage.

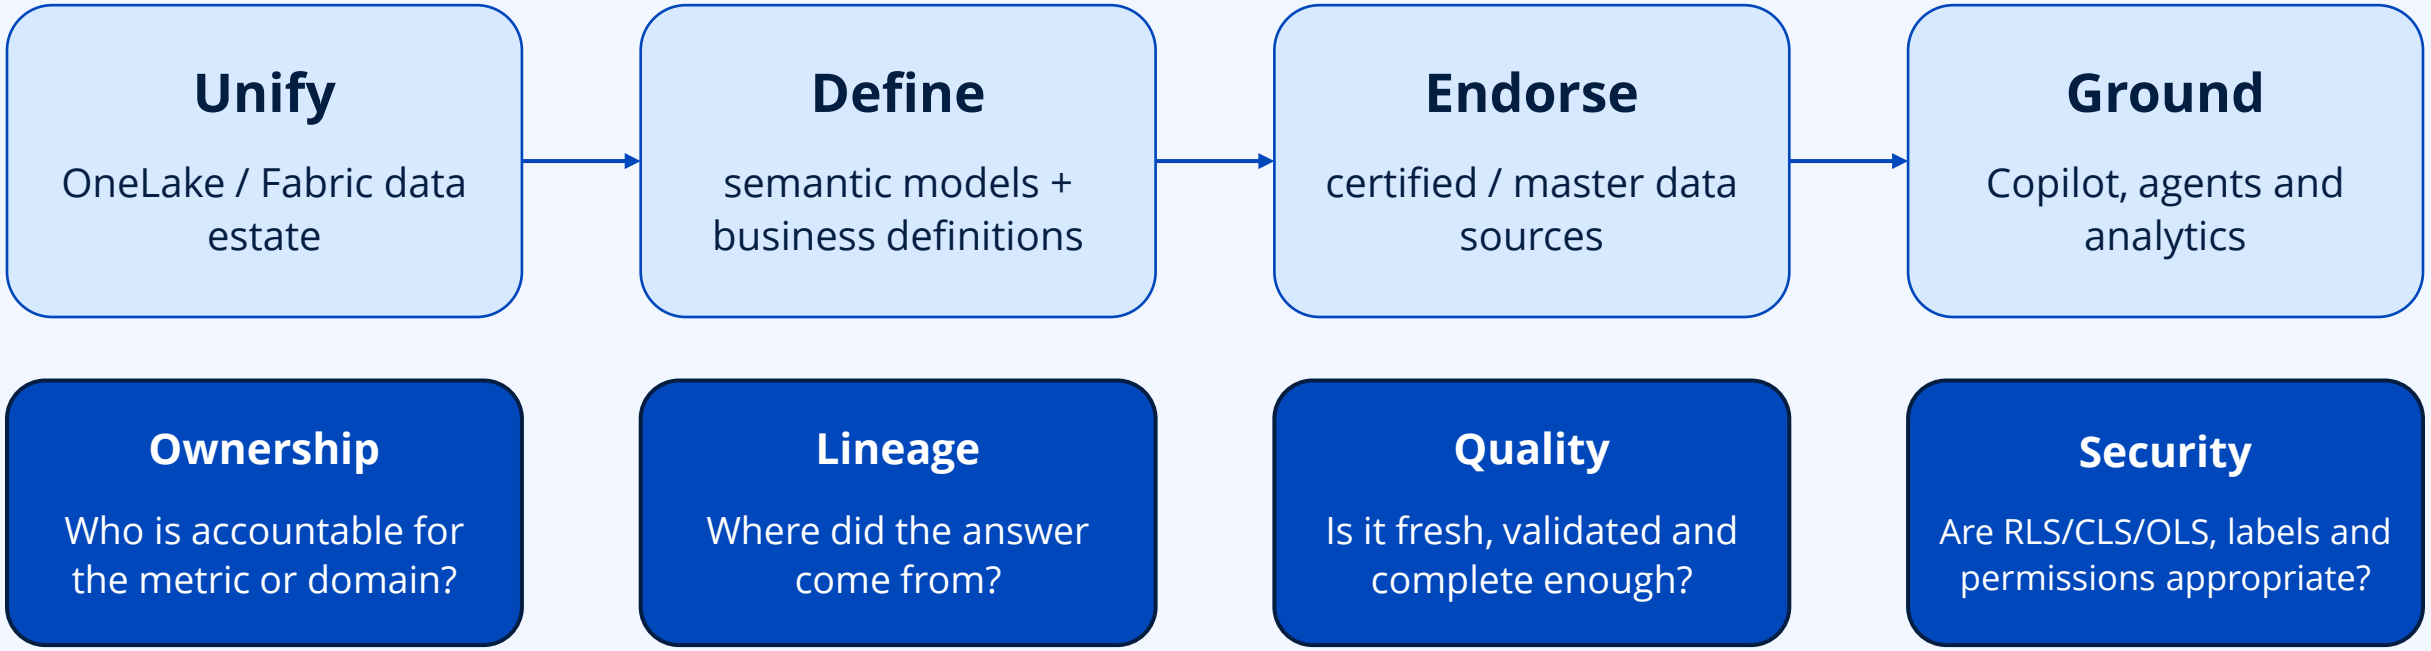

Agents turn answers into delegated action

Once AI can call tools, send messages or update systems, governance must cover authority to act, not just retrieval.

Example agent path Read SharePoint → query Fabric → draft recommendation → email customer team → update a workflow

Treat agents as delegated authority

A chatbot setting is too weak. Agents need identity, lifecycle, tool governance, approvals and auditability.

Identity

Which identity does the agent use?

Data

Which sources may it retrieve from?

Tools

Which connectors or systems can it call?

Approval

Which actions require human approval?

Evidence

Can we observe and investigate it?

Exit

Who owns, monitors and retires it?

Agent 365 control pattern

inventory • lifecycle • least-privilege tool access • sensitive-data protection • human approvals • audit trail • kill switch

Bottom line: an agent's permissions are delegated authority — not a productivity preference.

Control identity, entitlement, context, action and evidence

From 2 trusts to 6 control levers: This is the practical translation of Zero Trust for Copilot, Microsoft IQ and agents.

Identity

Who or what is asking?

Entra · MFA · Conditional Access

Entitlement

What may it reach now?

SharePoint · Teams · OneDrive · access reviews

Context

What should it trust?

Fabric IQ · semantic models · endorsement

Interaction

What may AI process or produce?

Purview labels · DLP · retention

Action

What may an agent do?

Agent 365 · connectors · approval gates

Evidence

Can we explain and revoke?

Audit · eDiscovery · Defender · lifecycle

A single AI answer may touch every control surface. The operating model must connect them.

Move from exposure to trusted AI context

The goal is not a perfect estate. It is a verified path from one priority business scenario to governed context and safe action.

0–30 days

Discover exposure + context

- Inventory Copilot/agent use cases
- Find sensitive, broad or ownerless data
- Identify authoritative metrics and semantic models
- Review current agents and connectors

Deliverable: exposure inventory + context map

31–60 days

Govern + connect

- Remediate access and ownership
- Apply labels, DLP and secure defaults
- Certify or endorse authoritative data
- Define agent privileges and approval gates

Deliverable: protected context architecture

61–90 days

Activate + scale

- Pilot one trusted business scenario
- Measure quality, traceability and adoption
- Monitor sensitive interactions and agent actions
- Package expansion backlog

Deliverable: governed use case + scale plan

Checkpoint: do not expand a use case until required access, protection, authoritative context and monitoring are verified.

From a business question to trusted AI at scale

Start with the decision the business needs to make, then make sure AI has the right context, controls and authority to support it.

Start with the decision	Map the context	Close the trust gaps	Activate safely	Scale what works
• Pick one high-value business question • Define the outcome and decision owner	• Identify the work, data and knowledge AI needs • Determine which sources are authoritative	• Fix access, ownership, quality and protection gaps • Define agent/tool permissions and human approvals	• Pilot with approved grounding and controls • Measure answer quality, usefulness and risk	• Expand to more domains and agents • Monitor access, quality, drift and usage continuously

What you leave with: Prioritized use case • Trusted-context map • Governance gaps • Pilot plan • Scale criteria

Eight questions before you scale Copilot or agents

These questions expose whether you have a trusted-context architecture or only a productivity pilot.

1. Can we find high-risk data before users or agents do?
2. Can we suppress discovery without pretending it is security?
3. Can we stop sensitive prompts, sources or outputs where required?
4. Can we inventory agents, tools, owners and actions?

5. Can we explain why a user or agent had access?
6. Can we identify the authoritative answer when sources conflict?
7. Can we prove which sources grounded an answer?
8. Can we revoke access or action rights fast enough?

If the answer is “not yet,” keep the use case narrow and close the control gap before scaling.

The winning posture is trusted context, not locked-down data

Customers will move faster when they know what AI can access, what it should trust and what it can safely do.

1

Copilot readiness is context readiness

Access, classification, lifecycle, ownership and authoritative sources determine trust.

2

Permission is not proof of current need

AI exposes stale and excessive access faster than humans ever could.

3

Agents make governance continuous

Once AI can act, identity, tools, approvals and evidence become operational requirements.

A practical customer conversation starts with three questions:

What can AI reach? • Which sources should AI trust? • What actions may an agent take?

The goal is not to minimize the data AI can use. The goal is to maximize the trusted context AI can use safely.

It's LEGO® Icons
Giveaway Time!

datavail

Visit Datavail at **Booth 303** to build your own **LEGO® Minifigure** and get more chances to **win!**

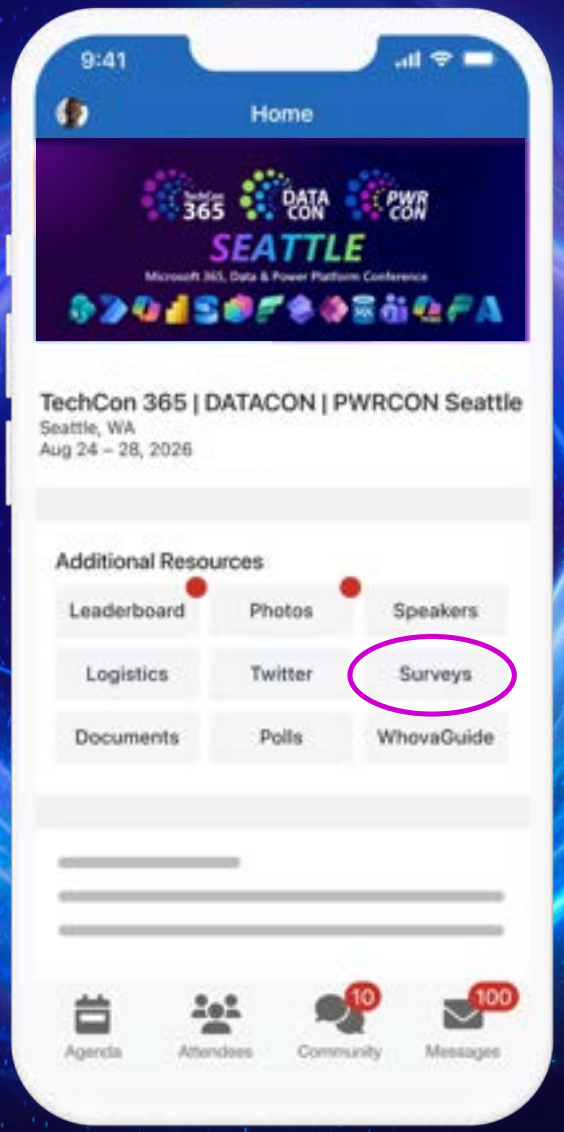

How was the session?

Search for Whova in the App Store or Google Play

Fill out the Session Surveys in the TechCon 365, DATACON, PWRCON Event App and be eligible to **WIN PRIZES!**