



# AI & Hyperscale Datacenter – Perspectives from Microsoft

Tec Forum – March 12 2024

Primo Amrein, Cloud Lead, Microsoft Switzerland



# Agenda

---

- **Microsoft's Sustainability Focus & AI**
- **Shared Responsibility and Customer Tools**
- **Hyperscale Datacenters & Sustainability**
- **Datacenter and Infrastructure Innovation**

# Microsoft Sustainability Priorities



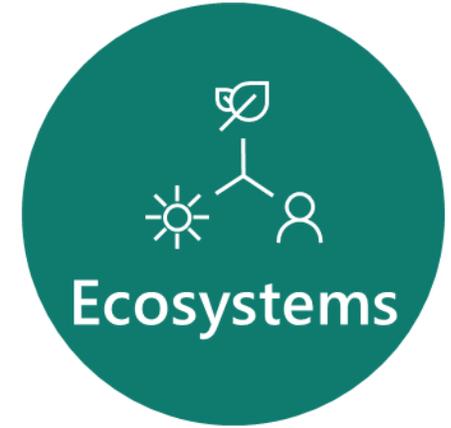
- Carbon negative by 2030
- Remove historical emissions by 2050
- \$1 billion climate innovation fund



- Zero waste by 2030
- Increase our reuse of servers and components up to 90% by 2025
- Invest \$30 million in circular economy



- Water positive by 2030
- Digitize water data
- Partner with Water.org and WRC members
- Invest \$10 million in water strategy fund



- Build and deploy a planetary computer
- Protect more land than we use by 2025
- Speak out on policy issues

Leadership Commitment

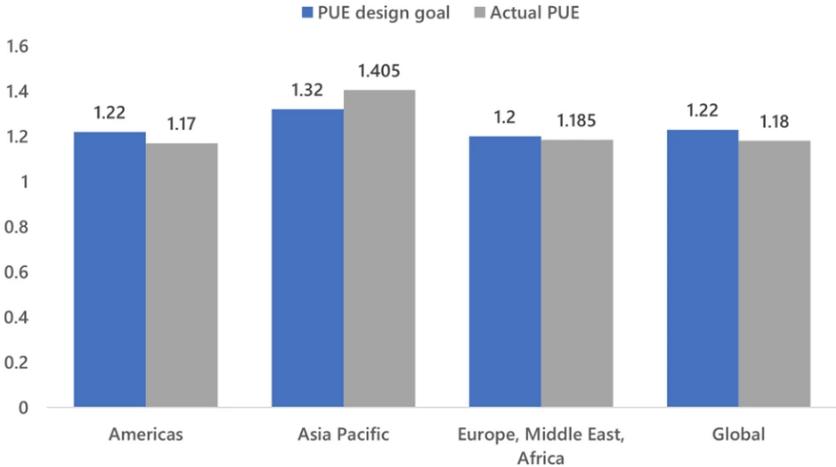
Accountability

Transparency

Advocacy

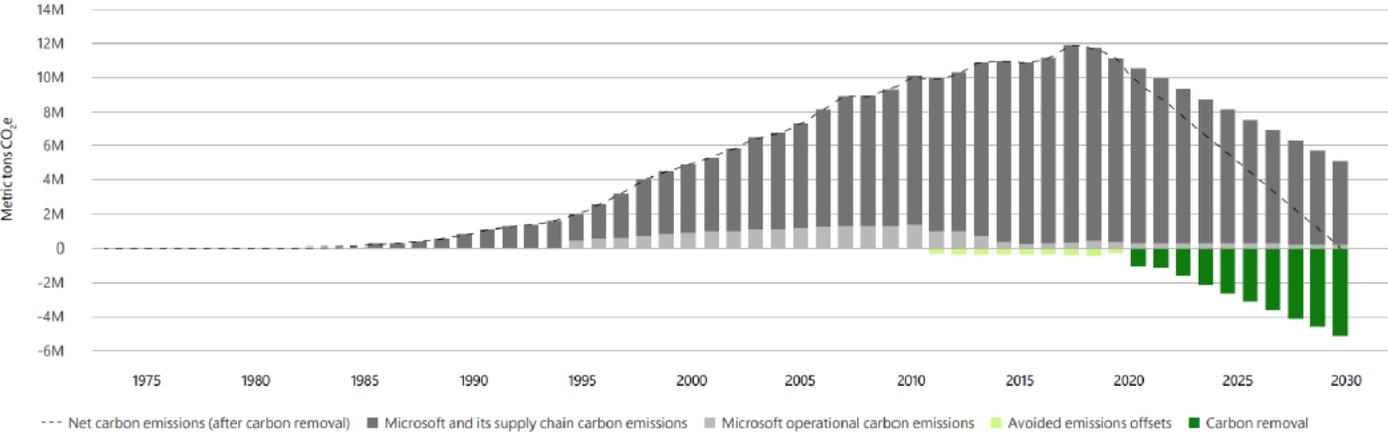
Co-innovation

# Microsoft's Pathway to Carbon Negative



Microsoft datacenter power usage effectiveness (PUE) =  $\frac{\text{total energy needed for facility}}{\text{total energy used for computing}}$

Microsoft's pathway to carbon negative by 2030<sup>1</sup>



<sup>1</sup> Chart has been updated to reflect latest actual values which incorporate latest methodology and structural change adjustments. Historic and projected emissions driven by latest actual data have also been updated for consistency. Projected removal values have also been updated consistent with our latest projected emissions and procurement strategy.

# AI Supporting Sustainability Efforts

## Using AI, Microsoft and UNFCCC take aim at emissions

An AI-powered platform will analyze carbon data for Paris Agreement signatories.

— MAEVE ALLSUP | NOVEMBER 30, 2023

[Link to Article](#)

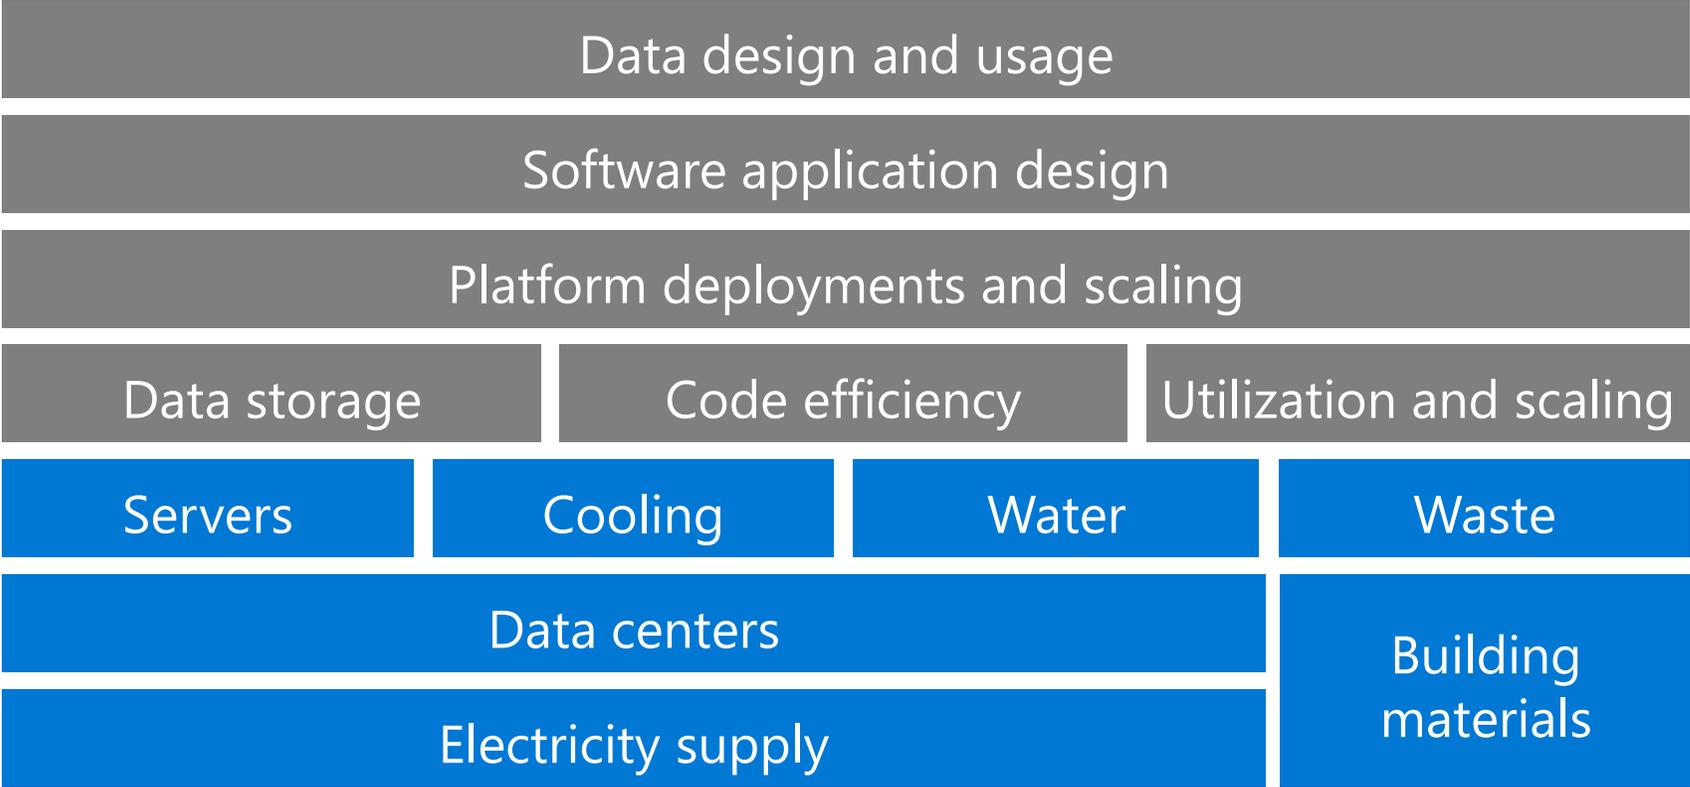
## Accelerating Sustainability with AI: A Playbook

Nov 16, 2023 | Brad Smith, Vice Chair and President, and Melanie Nakagawa, Chief Sustainability Officer

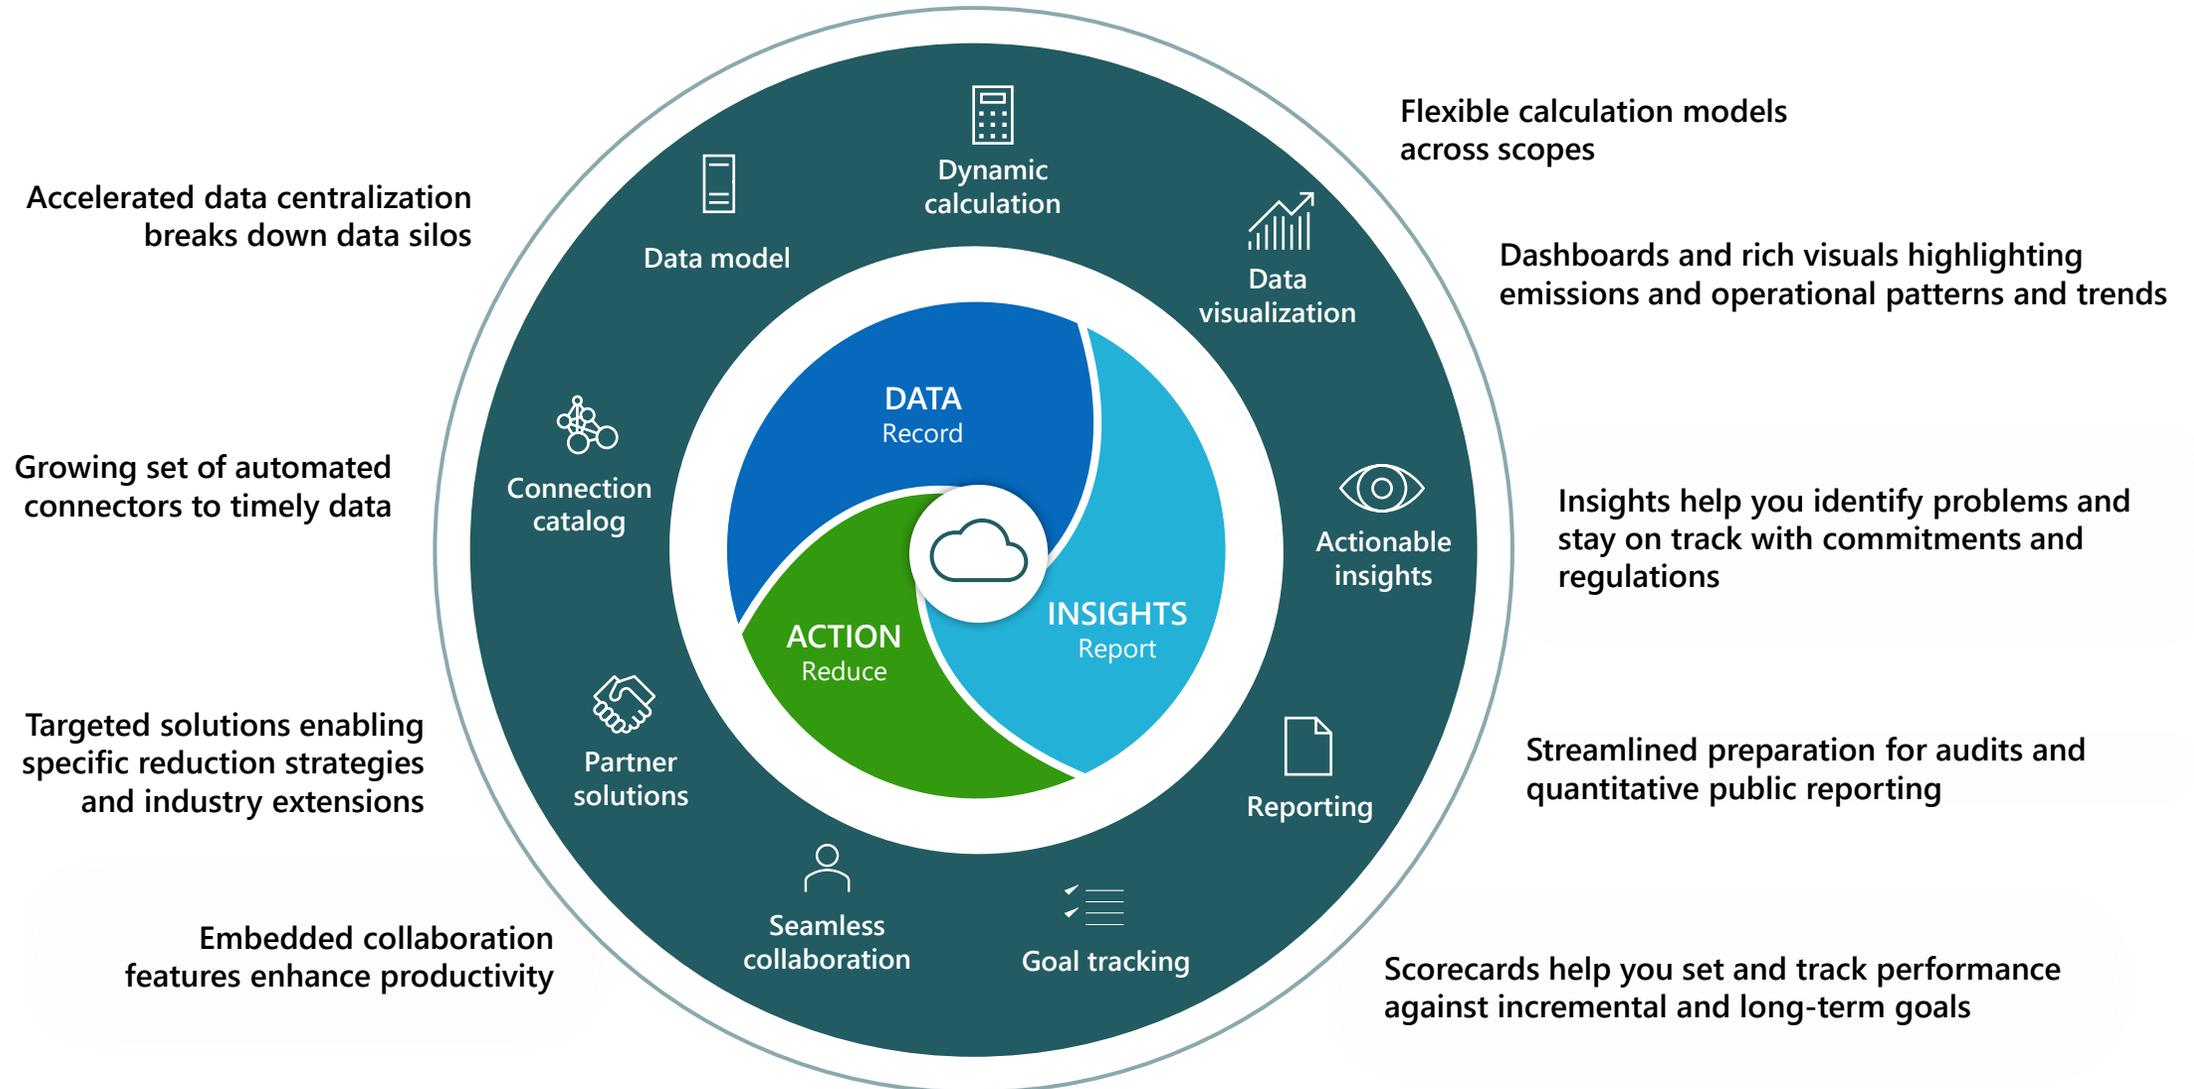


[Link to Article](#)

# Shared Sustainability Responsibility

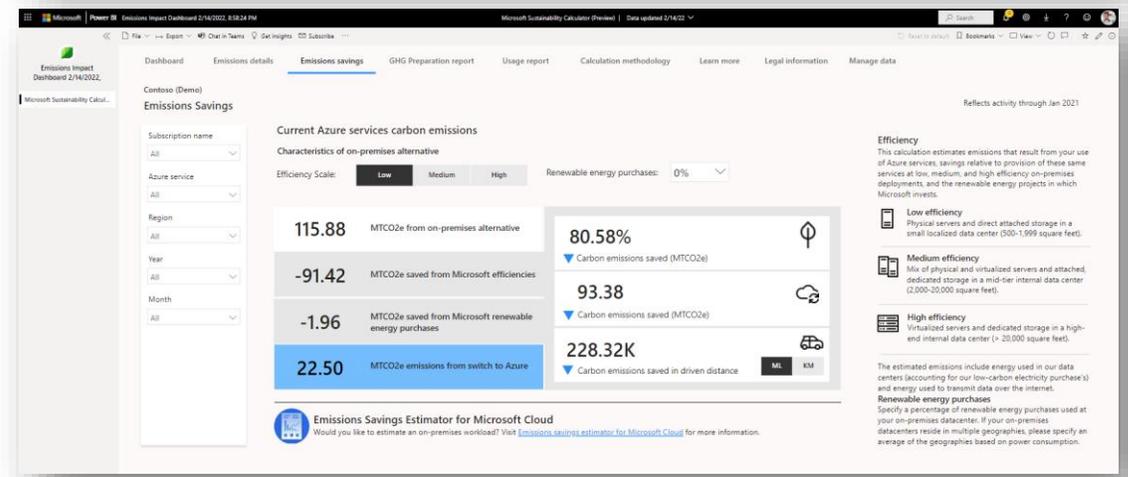
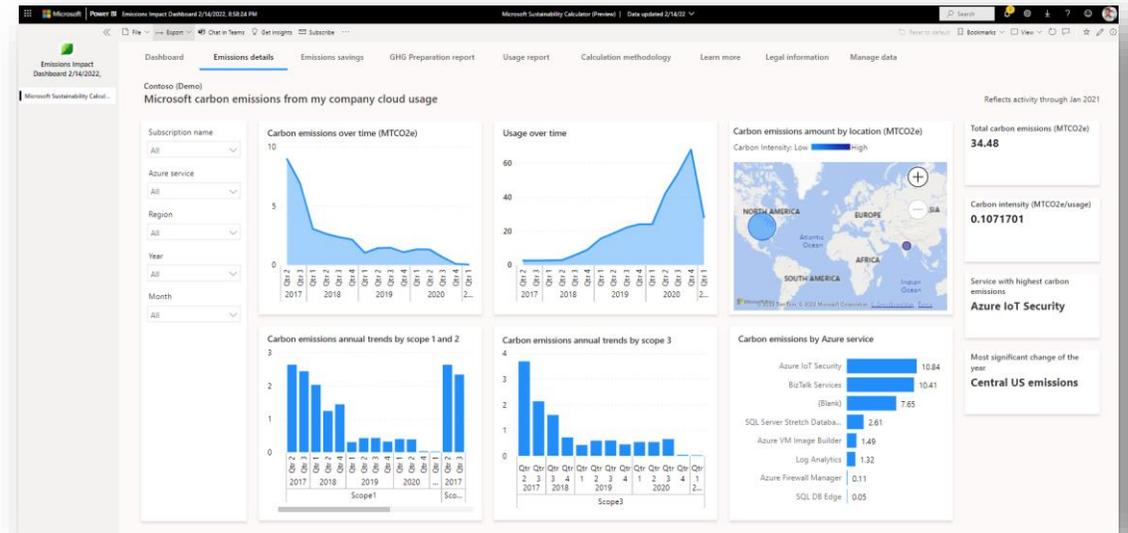


# Microsoft Sustainability Manager



# Emissions Impact Dashboard

Estimate your carbon emissions—and emissions you've already saved—related to using Microsoft cloud services



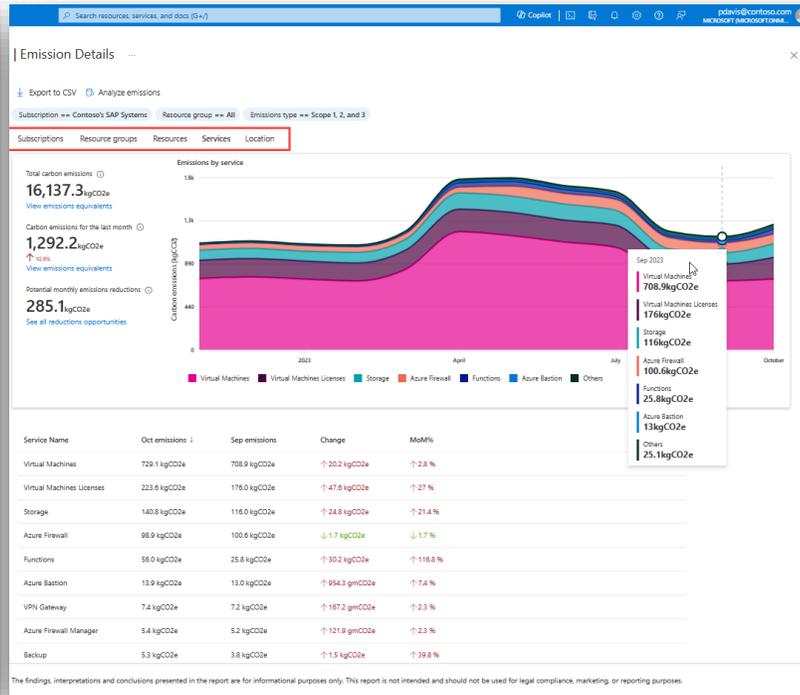
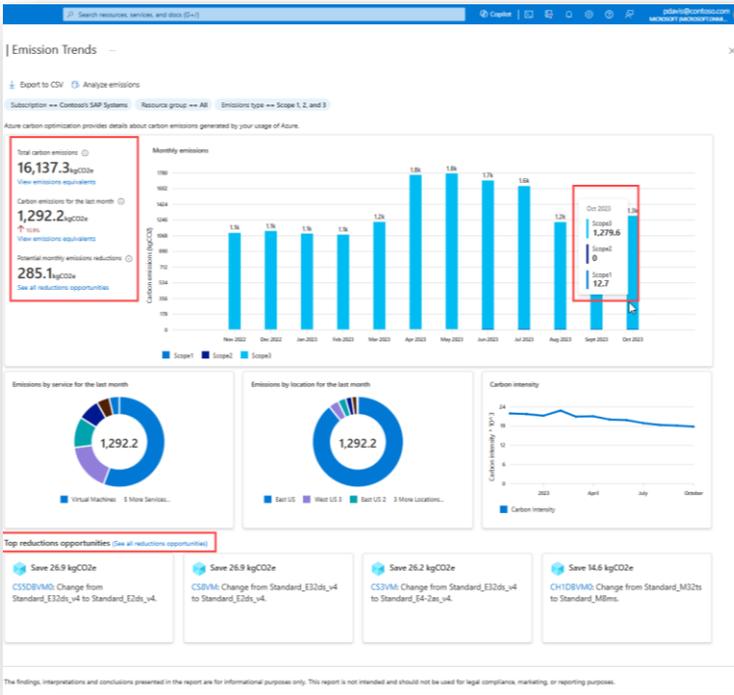
# Azure Carbon Optimization (preview)

Optimize your resource utilization to lower your carbon emissions and costs

## Emission Trends

## Emission Details

## Emission Reductions



**Total recommendations**  
55

**Potential monthly emissions reductions**  
285.1 kgCO<sub>2</sub>e

**Carbon reductions equivalent**  
5 planted trees more equivalents

**Potential monthly cost savings**  
\$3,663

| Subscription          | Resource group | Resource | Recommendation                                 | Emissions Reductions     | Cost Savings |
|-----------------------|----------------|----------|--|--------------------------|--------------|
| Contoso's SAP Systems | contoso-10     | CS5DBVM0 | Change from Standard_E32ds_v4 to Standard_E... | 26.9 kgCO <sub>2</sub> e | \$342.00     |
| Contoso's SAP Systems | contoso-13     | CS3VM    | Change from Standard_E32ds_v4 to Standard_E... | 26.9 kgCO <sub>2</sub> e | \$342.00     |
| Contoso's SAP Systems | contoso-5      | CS3VM    | Change from Standard_E32ds_v4 to Standard_E... | 26.2 kgCO <sub>2</sub> e | \$344.00     |
| Contoso's SAP Systems | contoso-6      | CH1DBVM0 | Change from Standard_M32ts to Standard_M8...   | 14.6 kgCO <sub>2</sub> e | \$202.00     |
| Contoso's SAP Systems | contoso-6      | CH1DBVM1 | Change from Standard_M32ts to Standard_M8...   | 14.6 kgCO <sub>2</sub> e | \$202.00     |
| Contoso's SAP Systems | R2-redhat-rg   | LT2DBVM1 | Change from Standard_M32ts to Standard_M8...   | 14.6 kgCO <sub>2</sub> e | \$202.00     |
| Contoso's SAP Systems | contoso-3      | CH3DBVMR | Change from Standard_E16ds_v4 to Standard_E... | 12.6 kgCO <sub>2</sub> e | \$160.00     |
| Contoso's SAP Systems | acss-demo-rg2  | DM2DBVM0 | Change from Standard_E16ds_v5 to Standard_E... | 12.6 kgCO <sub>2</sub> e | \$160.00     |
| Contoso's SAP Systems | acss-demo-rg2  | DM2DBVM1 | Change from Standard_E16ds_v5 to Standard_E... | 12.6 kgCO <sub>2</sub> e | \$160.00     |
| Contoso's SAP Systems | acss-demo      | D01DBVM1 | Change from Standard_E16ds_v4 to Standard_...  | 11.5 kgCO <sub>2</sub> e | \$143.00     |

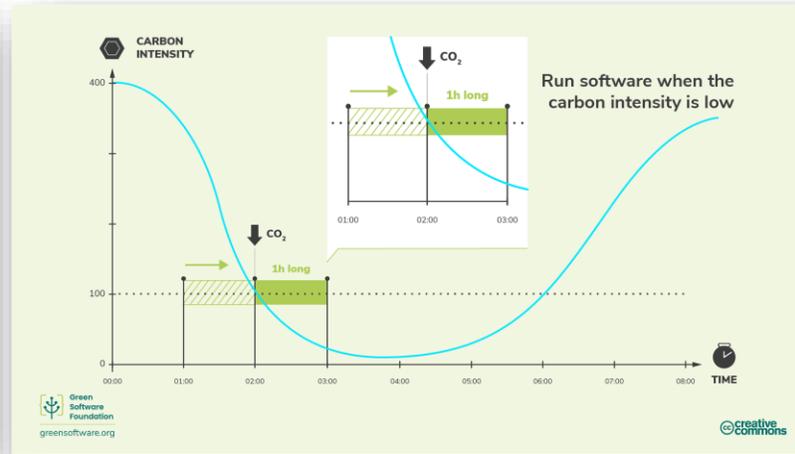
# Carbon Aware Computing

Do more when energy comes from low carbon sources and do less when energy comes from high carbon sources

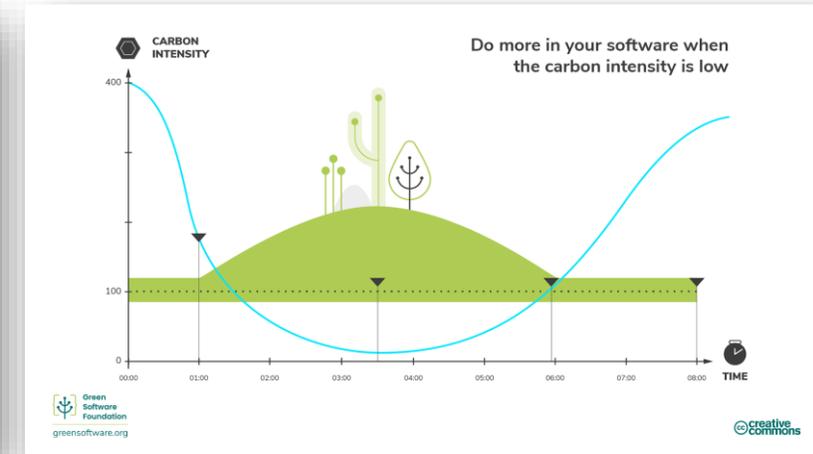
Demand Shifting: Spatial



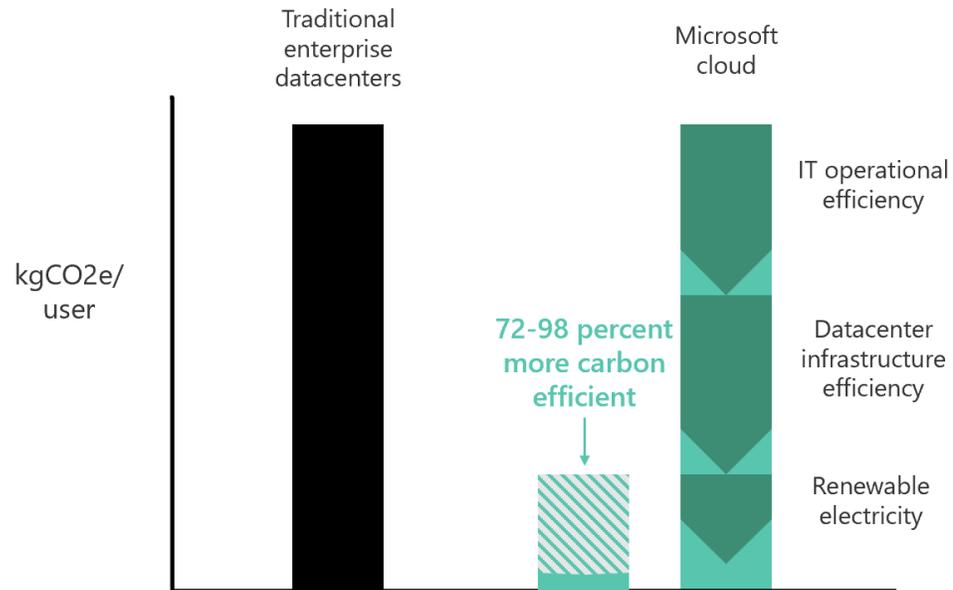
Demand Shifting: Temporal



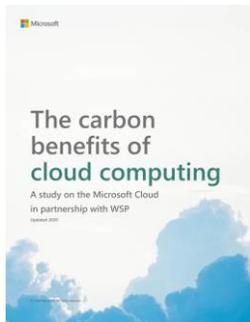
Demand Shaping



# Better Sustainability thanks to Hyperscale Datacenters



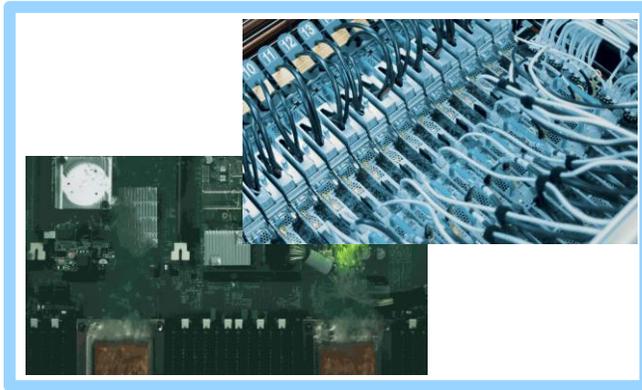
- Very high economies of scale result in much greater IT operational efficiency than smaller data centers
- More specialized, more efficient IT equipment alone reduces power consumption by 10 percent or more
- Advanced infrastructure technologies in hyperscale data centers reduce energy requirements for common tasks such as lighting, cooling, and power conditioning
- Large-scale green electricity purchases and renewable energy projects due to high consolidated demand



[www.microsoft.com/sustainability](http://www.microsoft.com/sustainability)

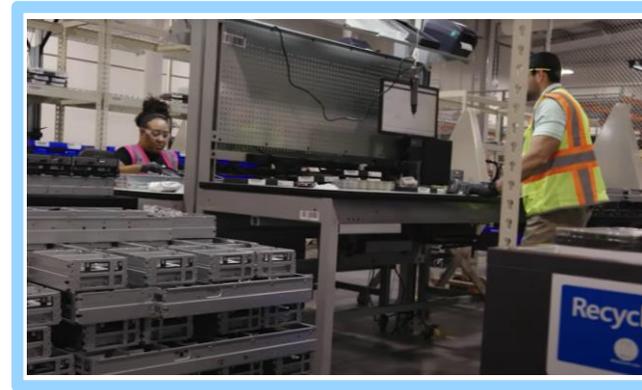
[Link](#) to the study

# Microsoft Innovation in Hyperscale Datacenters



## Energy efficiency

Reduced energy consumption thanks to liquid cooling



## Circular Economy

Circular Centers recycle IT-components onsite



Microsoft datacenter batteries to support growth of renewables on the power grid

## Renewables storage

DC batteries temporarily store renewable energy



## Hydrogen Fuel Cells

R&D for more sustainable backup power needs

# Azure AI Infrastructure

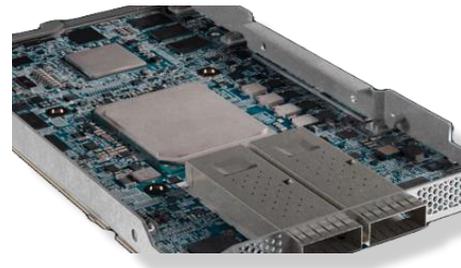
AI-optimized hyperscale infrastructure



## Network

Hollow Core Fiber (HCF)

**47x** Improvement in networking speed



## Performance

**Azure Boost:** Offloads server virtualization, freeing up CPU resources for improved performance

- **Networking speeds up to 200 GBps** networking speeds
- **Remote storage up to 12.5 GBps** throughput 650K IOPS
- **Local storage up to 17.3 GBps** throughput 3.8M IOPS

## Chip to cloud AI innovation



**Azure Cobalt** (CPU)  
Fastest of any cloud provider



**Azure Maia** (AI accelerator)  
More efficient LLM training & inference



**Azure NCCv5\*** (GPU)  
Azure Confidential GPU VMs for AI.  
Codesigned by Microsoft + Nvidia

**ND H200 v5\*** (GPU)  
Larger models and faster inferences



**AMD MI300X\*** (AI accelerator)  
for AI optimized Azure VMs

\*Nvidia ND H200 v5 is coming soon. Azure NCCv5 and AMD MI300X are in preview and early access, respectively.

# AI Hardware Innovations



A custom-built rack for the Maia 100 AI Accelerator and its "sidekick" inside a thermal chamber at a Microsoft lab in Redmond, Washington. The sidekick acts like a car radiator, cycling liquid to and from the rack to cool the chips as they handle the computational demands of AI workloads. Photo by John Brecher for Microsoft.