Data-Intensive Applications

Analytics that scales **efficiently** to search graphs at **extreme scale**

Applications of graph theory abound – spanning disciplines as varied as mathematics and the physical sciences, to the life and social sciences; and, graph analytics underpins numerous, popular algorithms in Machine Learning. By enabling Top30-rank processing capabilities of Graph500-scale problems, Microsoft and AMD empower stakeholders with critical insights obtained efficiently.

Successful data-intensive applications demand attention to several **business challenges:**

Search Efficiency

Rapidly traverse edges to find and label vertices for lightning-fast search.

Data-Driven Analytics

Effectively manipulate 'terascale' volumes of data (TBs) in concert with graph-analytics computations.

Big Data 6Vs

Beyond volume, fully account for the variety, velocity, veracity, validity and volatility of your data.

Graph analytics exists at the very core of a multitude of applications spanning several disciplines. To deliver search that scales, data-intensive applications demand support for distributed-memory parallel computing on Microsoft Azure:

HBv2 VM

Top 28 effective GTEPS ranking on BFS benchmark

GRAPH 500

Azure HBv2 Virtual Machines

HBv2 VMs feature 120 AMD EPYC 7002 Series CPU cores, 340 GB/s of memory bandwidth, and Mellanox 200 Gigabit/sec HDR InfiniBand™.

AMD EPYC™ 7002 Virtual Machines

The EPYC 7002 offers 45% higher memory bandwidth than competitive alternatives and PCIe Gen 4.0 to support the most-advanced networking.

Graph500 Breadth-First Search (BFS)

A BFS of a graph starts with a single source vertex, then, in phases, finds and labels its neighbors, then the neighbors of its neighbors, etc. Many graph algorithms make use of the BFS method.

Microsoft Azure

NDA Only

AMD

Platform performance benchmarks & validation data

Graph500 Breadth-First Search Method

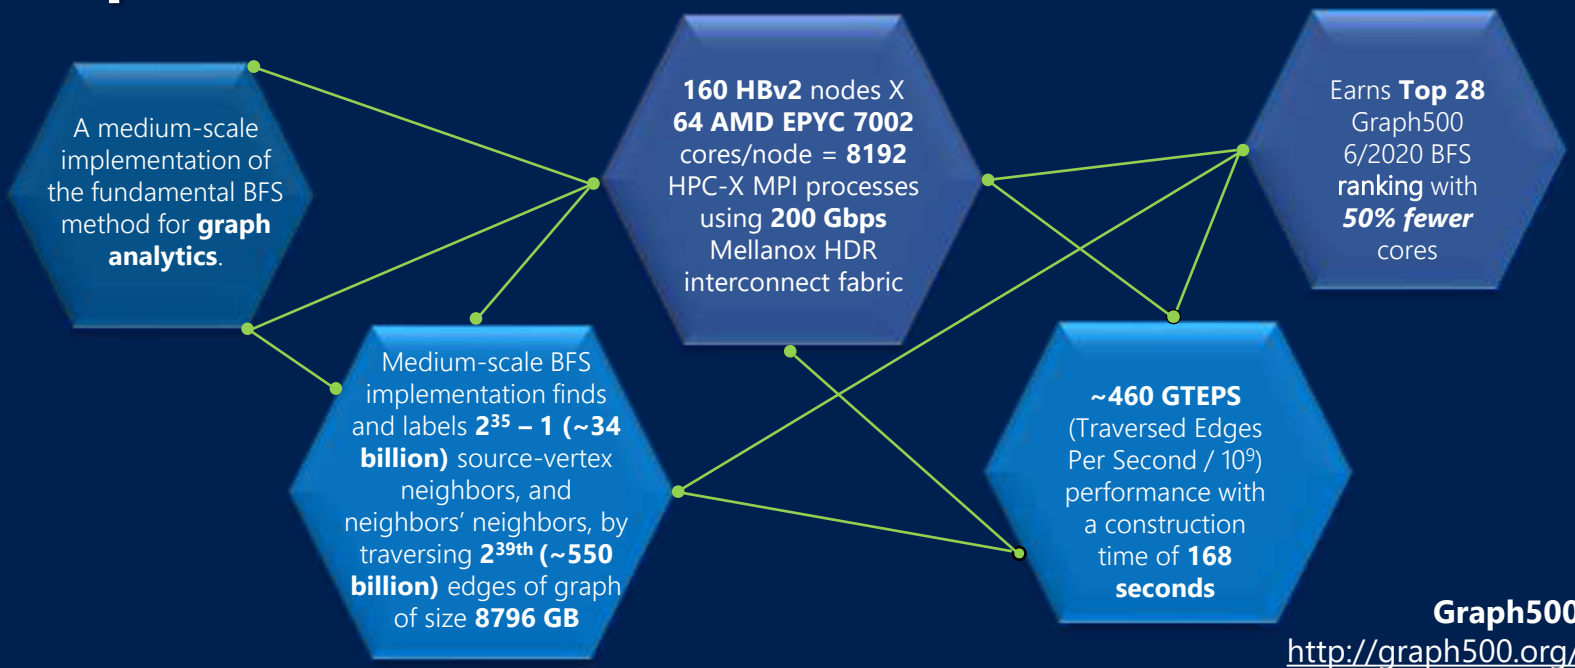

Graph500
http://graph500.org/

Microsoft Azure

HBv2

Virtual Machines

Azure HBv2 virtual machines feature 120 EPYC 7002 Series Processors from AMD. These VMs offer supercomputer-class performance, MPI scalability, and cost efficiency for a variety of real-world high-performance computing (HPC) and AI workloads, such as CFD, explicit finite element analysis (FEA), graph analytics, reservoir modeling, rendering, seismic processing, and weather simulation. Specifications:

CPU cores	Memory	Memory per core
120	480 GB	4 GB

Local SSD: GiB	RDMA network	Azure network
1.6 TB	200 Gbps	40 Gbps

AMD

EPYC 7002

Series Processors

AMD EPYC 7002 Series Processors unlock performance and redefine economics for HPC on Azure. AMD works with the open source community to help ensure your applications work exceptionally well with EPYC. AMD's comprehensive coverage of software compatibility and certifications are why Microsoft Azure trusts AMD EPYC processors for its most demanding services. AMD EPYC enables Azure HBv2 customers to achieve ground-breaking HPC performance at a competitive price point.

45%
more memory bandwidth than competitive alternatives

PCIe 4.0
supporting advanced networking capabilities for tightly coupled workloads

Next Steps **Graph500 BFS Method on Azure**
(URL: https://github.com/jithinjosepkl/graph500)