Presentation Brief

RISC-V: Enabling a New Era of Open Data-Centric Computing Architectures

Delivers Independent Resource Scaling, Open Source, and Modular Chip Design for Big Data and Fast Data Environments

Prepared for:

7th RISC-V Workshop

November 28–30, 2017

Western Digital Corporation | Milpitas Facility

Western Digital®

Contents

SAFE HARBOR – DISCLAIMERS

Forward-looking Statements

This document contains forward-looking statements that involve risks and uncertainties, including, but not limited to, statements regarding our product and technology positioning, the anticipated benefits of our new technologies and transitioning into RISV-V open instruction set architectures. Forward-looking statements should not be read as a guarantee of future performance or results, and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in or suggested by the forward-looking statements.

Additional key risks and uncertainties include the impact of continued uncertainty and volatility in global economic conditions; actions by competitors; difficulties associated with go-to-market capabilities and transitioning into RISC-V open instruction set architectures; business conditions; growth in our markets; and pricing trends and fluctuations in average selling prices. More information about the other risks and uncertainties that could affect our business are listed in our filings with the Securities and Exchange Commission (the "SEC") and available on the SEC's website at www.sec.gov, to which your attention is directed. We do not undertake any obligation to publicly update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as otherwise required by law.

"The role of data is changing as well, evolving from just being a record or a log of events, recordings or measurements, to forms of communications that set the foundation for driving efficiencies in productivity and automation."

The Data-Centric World

Most all discussions about today's data begin with it growing at an exponential rate, doubling every two years that will reach hundreds of zettabytes next decade. This data deluge is driving the multiple Vs (volume, velocity, variety, and value), and requires a longer shelf life for future analysis, extracting further value and intelligence that enables better business and operational decisions.

What's driving this data-centric world is that the role of data is changing, evolving from just being a record or a log of events, recordings or measurements, to forms of communication that deliver efficiencies in productivity and automation, and ultimately, the value that data delivers, become a form of currency.

Also driving data growth is that the abundance of sources for data generation are increasing as well. Data no longer is just generated from applications, but it now comes from mobile devices, production equipment, machine sensors, video surveillance systems, Internet of Things (IoT) and industrial IoT (IIoT) devices, healthcare monitors and wearables, to name a few. And, the data that is being generated is created in both large-scale data centers at the "core," and in remote and mobile sources at the "edge" of the network.

With this paradigm shift, Big Data and Fast Data environments are spawning rapidly. Big Data represents very large data sets that may need to be analyzed using computations and algorithms, unmasking trends, patterns, and associations between seemingly disparate data sets. The analysis provides valuable insights and makes new connections that drive precise predictions and decisions that can help to achieve better outcomes. However, this data is based on information from the past. Today's applications need to respond to events as they happen.

Fast Data environments process or transform data as it is captured, leveraging the algorithms derived from Big Data to provide real-time decisions and results. As Big Data provides insights derived from 'what happened' to 'what will likely happen' (predictive analysis), Fast Data delivers real-time actions. The data for these environments can come from a host of 'smart' machines, environmental monitoring, security and surveillance systems, and securities trading systems that need to act in real-time.

Big Data and Fast Data form a virtuous cycle, as Big Data apps deliver more and more algorithms that enable Fast Data environments, which in turn, encourages more sources of data to be fed into Fast Data applications. The more data generated from Fast Data applications help to fuel more and larger Big Data environments, and thus, the cycle repeats.

"As Big Data gets bigger and faster, and Fast Data gets faster and bigger, the 'one size fits all' approach of general-purpose computing is failing to meet the increasingly diverse application workloads of our data-centric world."

Diversity of Applications and Workloads

As Big Data and Fast Data environments proliferate, they break the boundaries of traditional infrastructure and system architecture. The "general-purpose" technologies and architectures that have served the industry so well in the past, are reaching their limits of scalability, performance and efficiency. General-purpose workloads that are supported by general-purpose architectures typically have a uniform ratio of processing resources, such as operating system (OS) processing, specialty offload processing, memory, storage, and interconnect. As Big Data gets bigger and faster, and Fast Data gets faster and bigger, the 'one size fits all' approach of general-purpose computing is failing to meet the increasingly diverse application workloads of our data-centric world.

What today's data-centric architectures need are capabilities that enable more control over the blend of resources required. As the diversity of Big Data and Fast Data workloads expand, data-centric compute architectures will need the ability to scale resources independent of one another. The architectures for tomorrow will need to go beyond the limited resource ratios of general-purpose compute architectures and enable purpose-built solutions with optimized levels of OS processing, specialty processing, memory, storage, and interconnect. The extreme data and compute workloads for analytics, machine learning, artificial intelligence and smart systems demand purpose-built architectures.

Data-Centric Data Center Architecture

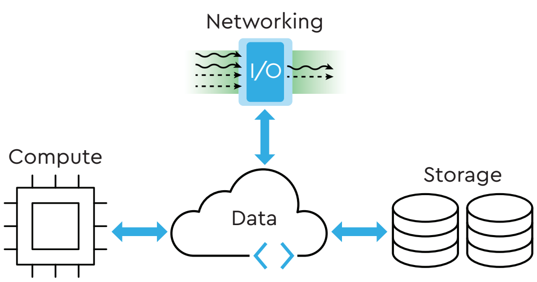

In addition to independent scaling of resources, data-centric environments at the edge bring further challenges. The remote and mobile conditions of the edge also require optimization of power consumption, environmental ruggedness, physical size, and weight. General-purpose compute architectures generally carry significant overhead in cost, size and power consumption, as they integrate general purpose logic to connect to peripherals, communications and input/output (I/O) devices that are not required for a purpose built environment.

Data-Centric Environments

With the increase in Big Data and Fast Data workloads, more purpose-built environments are required where resources can be independently scaled. One example includes Big Data analytical environments that need massive storage, but only moderate processing and interconnect since the results are not required in real-time. Machine learning environments are another example where a lot of storage is needed, as well as a lot of specialty processing (through GPUs, ASICs, and FPGAs), to accelerate learning. Another example includes the Fast Data application of real-time video event detection where facial recognition from a video feed is compared to a library of facial signatures to trigger a security alert, and requires large resource pools of specialty processors and memory.

> "As Big Data and Fast Data applications start to create more extreme workloads, purpose-built architectures will be required to pick up where today's general-purpose architectures have reached their limit."

To drive and achieve independent scaling, the following is required:

- An open architecture to avoid dependency on a particular vendor or technology as openness fosters accelerated development and adoption of a wide range of technologies allowing market applicability rather than proprietary barriers to determine deployment.
- Modular technologies and dense integration for remote or mobile edge environments that can optimize space, weight, cost, and power consumption, as building components with only the necessary functions and logic is what is needed. This approach eliminates the potential inefficiencies of general-purpose components, such as CPUs, that integrate unnecessary functionality, cost, size and power consumption.

As data-centric environments, and their needs, have been identified, the next step is to present how the RISC-V open architecture can be the foundation for enabling purpose-built data-centric computing environments.

Why RISC-V?

An instruction set architecture (ISA) is the set of machine instructions that comprises the machine language and I/O model of a computer or computing system. It literally defines everything that a programmer needs to know in order to program it.

RISC-V is an open, free ISA enabling a new era of processor innovation through open standard collaboration. It delivers a new level of free, extensible software and hardware freedom on architecture, paving the way for the next 50 years of computing design and innovation.

In contrast to most ISAs, RISC-V can be freely used for any purpose, permitting anyone to design, manufacture and sell RISC-V chips and software. RISC-V and any CPU design change are freely available under a Berkeley Software Distribution (BSD) license. Converse, commercial chip vendors charge license fees for the use of their proprietary source code and patents.

Based on its open, modular approach, RISC-V is ideally suited to serve as the foundation of data-centric compute architectures. As an OS processor, it can enable purpose-built architectures by supporting the independent scaling of resources. Its modular design approach allows for more efficient processors for edge and mobile systems.
As Big Data and Fast Data applications start to create more extreme workloads, purpose-built architectures will be required to pick up where today's general-purpose architectures have reached their limit.

The key benefits of RISC-V for enabling a new era of data-centric computing architectures include:

1. Independent scaling of resources
2. Open source capabilities
3. Modular chip designs

"This concept of bringing compute to data, rather than data to compute, creates more intelligent infrastructures and alleviates the need for costly increases in interconnect bandwidth as data repositories grow beyond the petabyte-scale level."

Commitment to RISC-V

Western Digital creates environments for data to thrive, whether supporting Big Data or Fast Data environments at the core, or at the edge. We expect that the open and scalable architectures enabled by RISC-V will help accelerate the deployment of data-centric applications in Big Data and Fast Data at the core, and at the edge.

With Western Digital transforming from a data storage company to a data technology company, expansion to open compute architectures is underway to bring processing closer to data. In today's 7th RISC-V Workshop, the company announced its commitment to help lead the advancement of data-centric compute architectures through the work of the RISC-V Foundation.

Moving Compute Closer to Data

When it comes to the capture, preservation, access, and transformation of data, the independent scaling of resources enabled by RISC-V can be applied at a number of levels – either across the entire data center, within a data center rack, or down to the system, platform, or device level. The ability to create purpose-built environments can be impactful at the macro-level, micro-level, and everywhere in-between. With an objective of surrounding data with the required processing, memory, storage, and interconnect resources, can be delivered in a number of ways.

For Big Data environments, the open nature of RISC-V fosters the ability to bring increasing amounts of processing to where the data resides. This concept of bringing compute to data, rather than data to compute, creates more intelligent infrastructures and alleviates the need for costly increases in interconnect bandwidth as data repositories grow beyond the petabyte-scale level. An implementation at the macro-level of infrastructure is likely to be in the embedding of data analytics or machine learning resources, and into high-scale data storage systems, such as Western Digital's ActiveScale™ multi-petabyte object-based storage (OBS) system.

At the micro-level of integration, application and specialized data processing can be added to storage devices, such as hard disk drives (HDDs) or solid state drives (SSDs), as well as storage platforms that integrate these devices into larger storage pools. This processing can be used to control the functionality of the storage devices and platforms, and provide onboard data processing and transformation capabilities.

There is also a third opportunity for supporting data-centric computing architectures at the edge. As mobile and edge-based Fast Data applications drive the need for the real-time transformation of data, RISC-V can enable more efficient solutions. The open and modular approach of RISC-V enables the right-sizing of resources, including the CPU itself and the tight integration of components, into a space- and power-efficient package.

"Western Digital is committed to the mission of the RISC-V Foundation and its ecosystem, and together, we can create environments for data to thrive."

Accelerating the RISC-V Ecosystem

As with any open source initiative, RISC-V will need a complete ecosystem surrounding it in order to thrive. Included in this ecosystem are open source contributors, test and validation groups, application development tools, chip foundries, system architects, solution developers and integrators, and users.

To justify the investment required to create the ecosystem, these groups will want assurances that the RISC-V initiative will ultimately be a significant source of technology and business value in the market going forward. To contribute toward the advancement and success of RISC-V, Western Digital brings a vested interest in the form of over one billion processor cores that it consumes on an annual basis across its product portfolio. The company is committed to advancing RISC-V technology for use in mission critical applications so that it can be deployed in its products. Western Digital is also engaged in active partnerships and investments in RISC-V ecosystem partners.

Looking Ahead

To realize the possibilities of data, we need to capture, preserve, access, and transform it to its full potential. Big Data and Fast Data environments in the core, and at the edge, have exceeded the capabilities of general-purpose compute architectures. The extreme data-centric environments of tomorrow's applications require purpose-built environments that support independent scaling of data transformation resources in an open manner.

Western Digital is committed to the mission of the RISC-V Foundation and its ecosystem, and together, we can create environments for data to thrive. When data thrives, our people, communities and planet can thrive, all through the power, potential and possibilities of data.

###

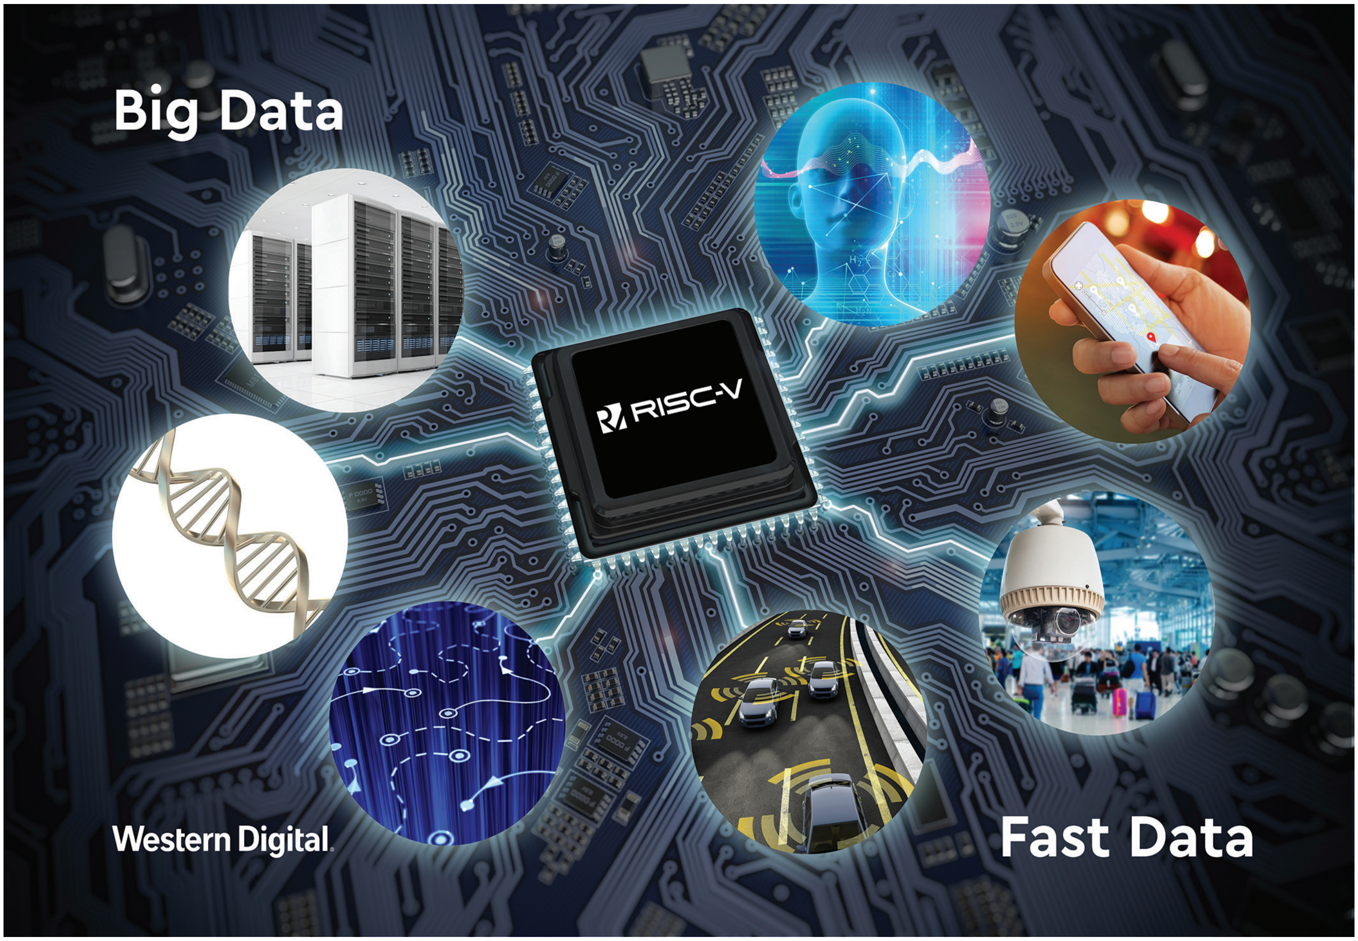

About Western Digital

Western Digital creates environments for data to thrive. The company is driving the innovation needed to help customers capture, preserve, access and transform an ever-increasing diversity of data. Everywhere data lives, from advanced data centers to mobile sensors to personal devices, our industry-leading solutions deliver the possibilities of data.

PRB-01-EN-US-1117-01