Architecting High-Speed Data Streaming Systems

Sujit Basu

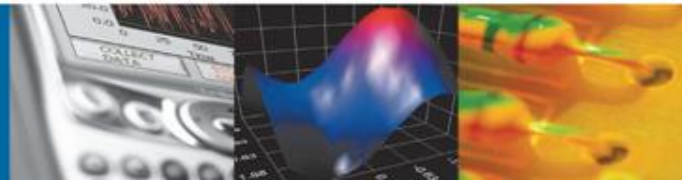

stream·ing [stree-ming] – verb

1. The act of transferring data to or from an instrument at a rate high enough to sustain continuous acquisition or generation.

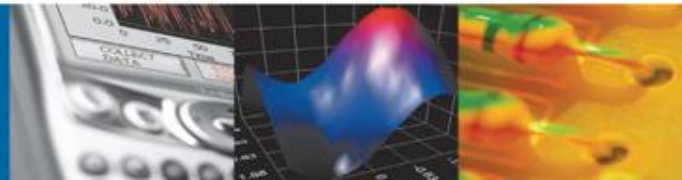

Motivation for Data Streaming

- Ever-increasing amounts of data
- Record “everything” and play it back later
- DAQ, DSA, vision, modular instruments
- Hard drives: faster, bigger, and cheaper
- RAID hardware is ubiquitous and inexpensive
  - PCI Express, ExpressCard, USB, 1394, eSATA
- PCI Express provides higher, dedicated bandwidth

Applications Requiring Data Streaming

- RF Recording and Playback
- IF and Baseband Streaming
- Noise Mapping
- Digital Streaming

Spectral Monitoring :
Chengdu Huari Telecommunications Company

RF Record and Playback:
B&B Technologies

Noise Mapping:
Boeing

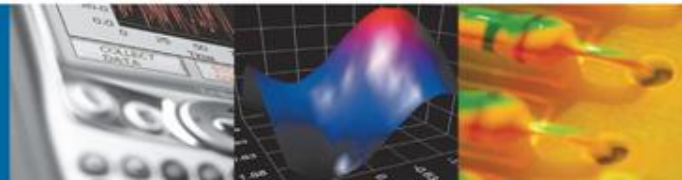

Key System Components

- **Hardware Platform with High-Throughput and Low-Latency**
- **High-Speed Data Storage**
  - Hard Drives (HDDs)
  - Solid-State Drives (SSDs)
- **Software for Streaming to Disk at High Rates**
- **Streaming Front-End Instrumentation**

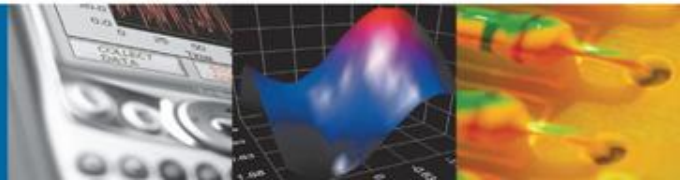

Key System Components

- **Hardware Platform with High-Throughput and Low-Latency**
- **High-Speed Data Storage**
  - Hard Drives (HDDs)
  - Solid-State Drives (SSDs)
- **Software for Streaming to Disk at High Rates**
- **Streaming Front-End Instrumentation**

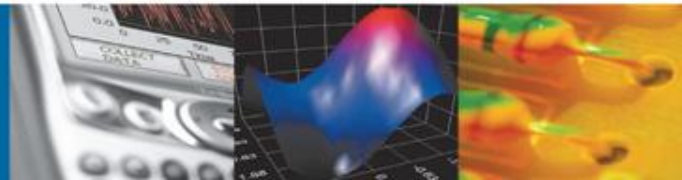

Bandwidth versus Latency

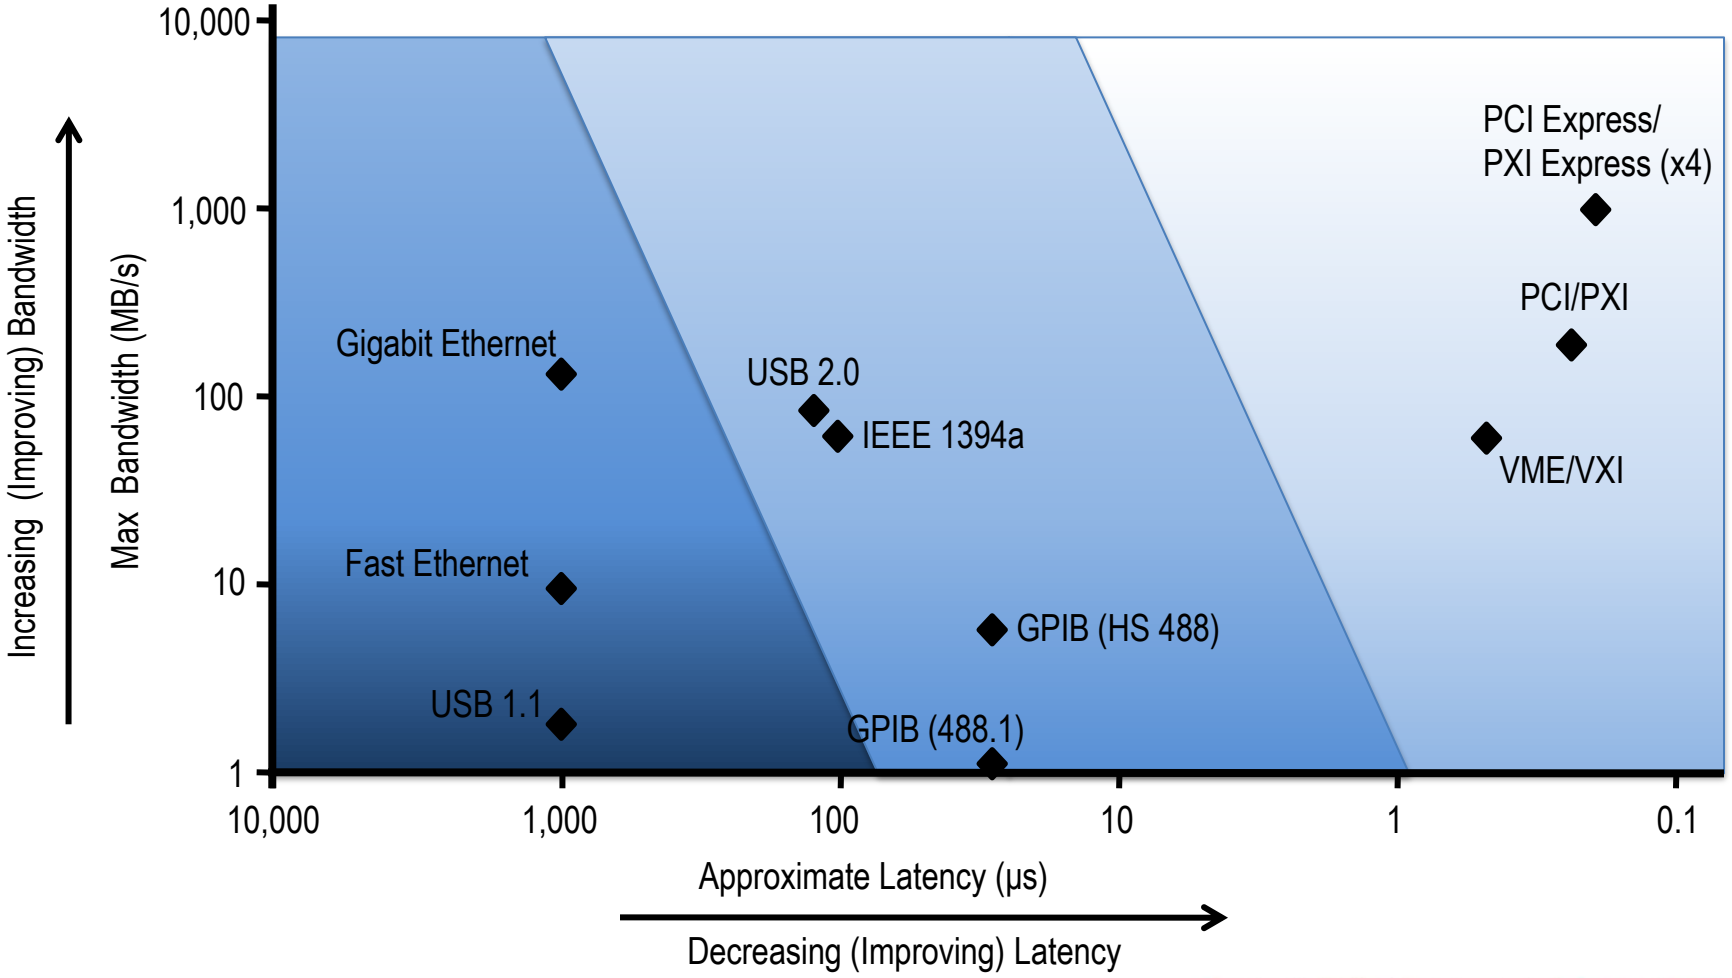

System Streaming Architecture – PXI

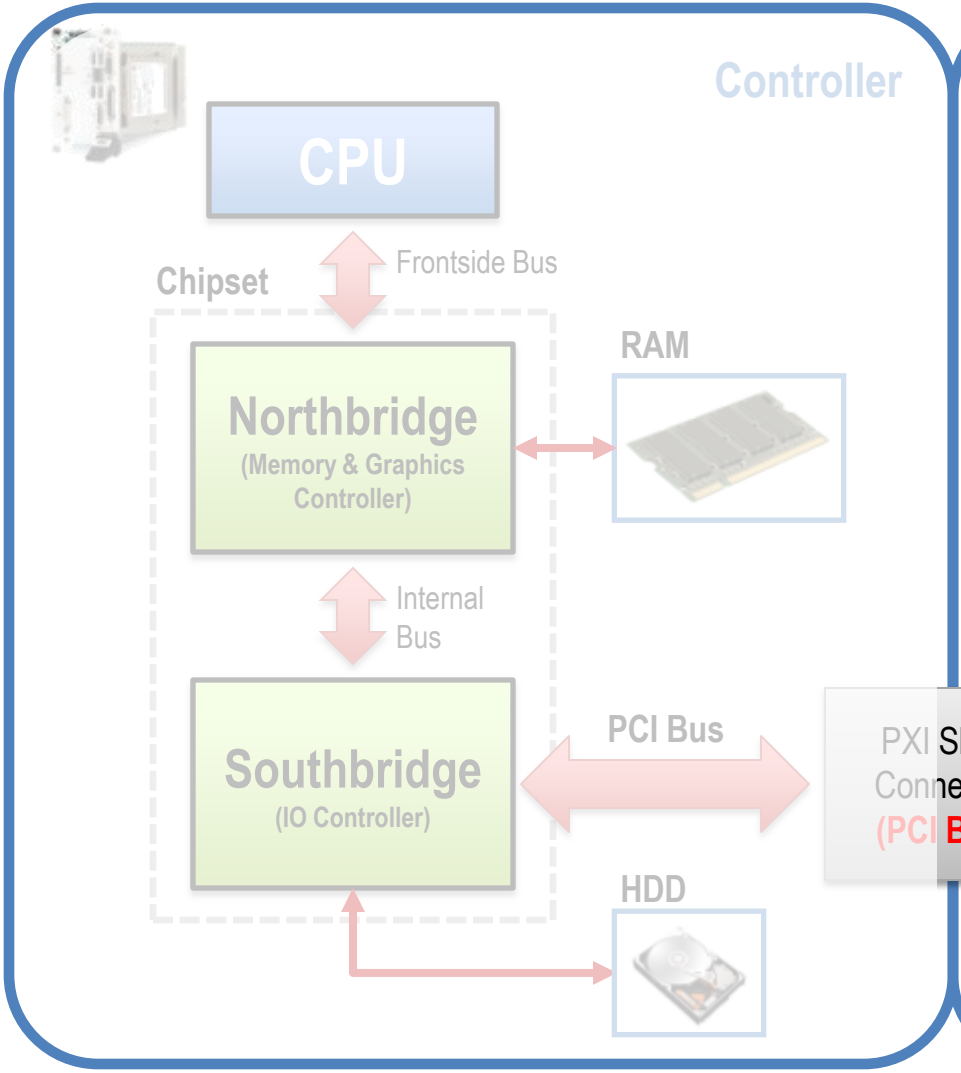

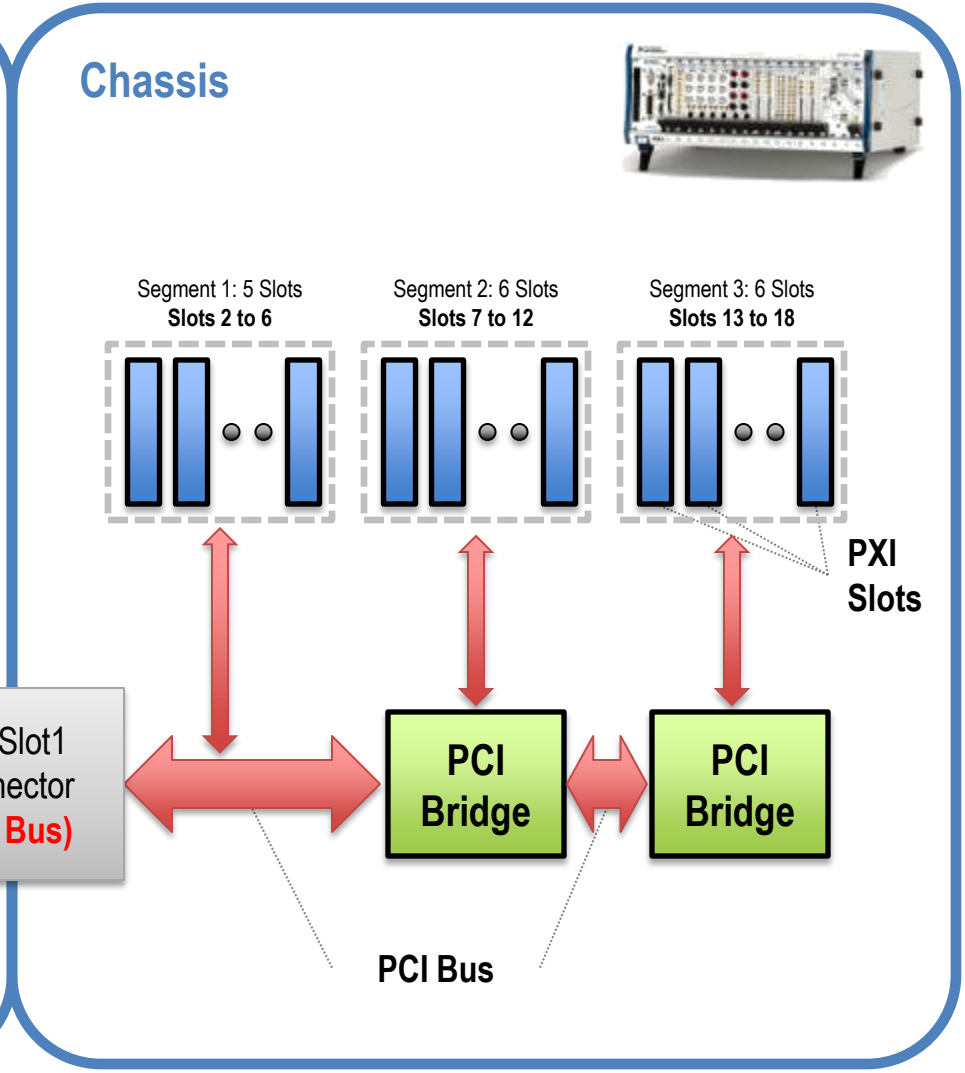

Streaming to/from Controller Hard Drives

Controller

CPU

Chipset

Frontside Bus

Northbri
(Memory and G ics
Controlle

RAM

rnal

Southbri e
(I/O Control

PCI Bus

PXI Slot1
(PCI Bus)

HDD

Chassis

Segment 1: 5 Slots
Slots 2 to 6

Segment 2: 6 Slots
Slots 7 to 12

Segment 3: 6 Slots
Slots 13 to 18

PXI
Slots

PCI
Bridge

CI
Bridge

PCI Bus

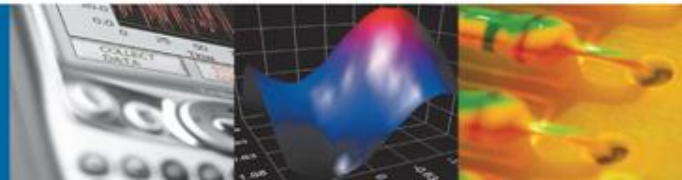

PCI Express Overview

- Serial interconnect at **2.5 Gbits/s**
  - PCI transactions are packetized and then serialized
  - Low-voltage differential signaling, point-to-point, 8 B/10 B encoded
  - x1 (by 1) has bandwidth of **250 Mbytes/s per direction**
  - x16 (by 16) has bandwidth of **4 Gbytes/s per direction**
- Uses same software model as PCI
  - Ensures software compatibility
- Road map for longevity with Gen 2 clocking (5 Gbits/s)

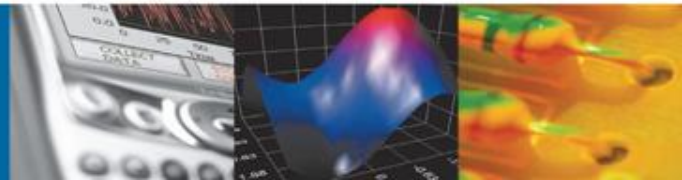

System Streaming Architecture – PXI Express

PXI Express System

NI PXIe-8130

NI PXIe-1062Q

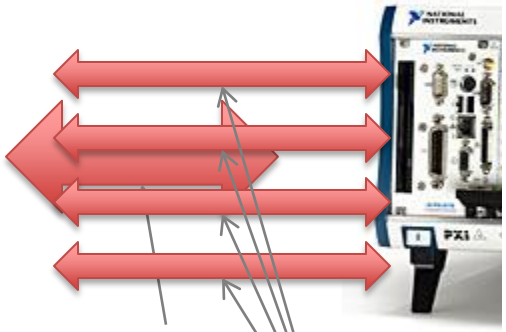

Four x4 PCI Express Links

Total system slot bandwidth depends on controller/chassis combination

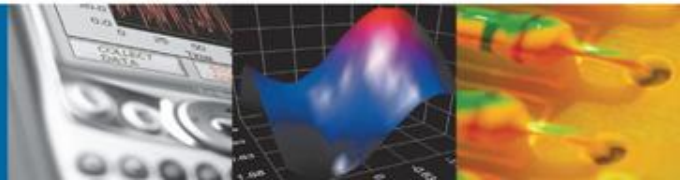

Streaming Architecture – PXI Express Controllers

NI PXIe-8130

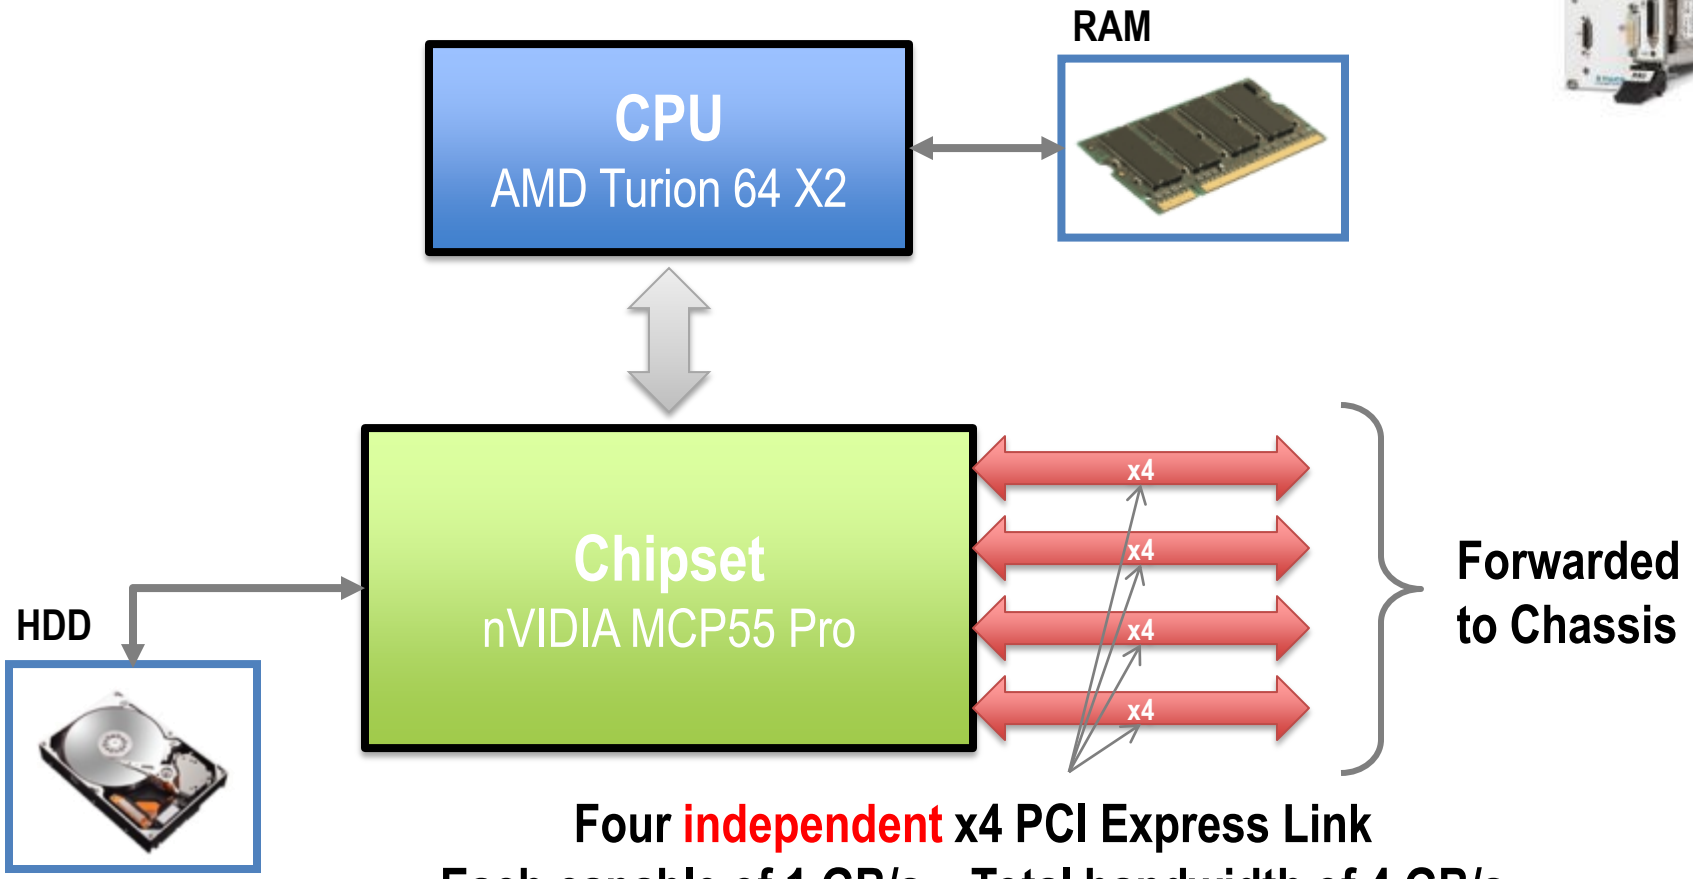

Four independent x4 PCI Express Link
Each capable of 1 GB/s – Total bandwidth of 4 GB/s

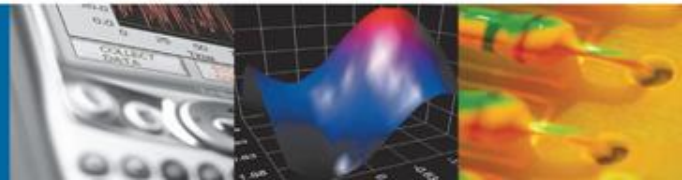

Streaming Architecture – PXI Express Controllers

NI PXIe-8108

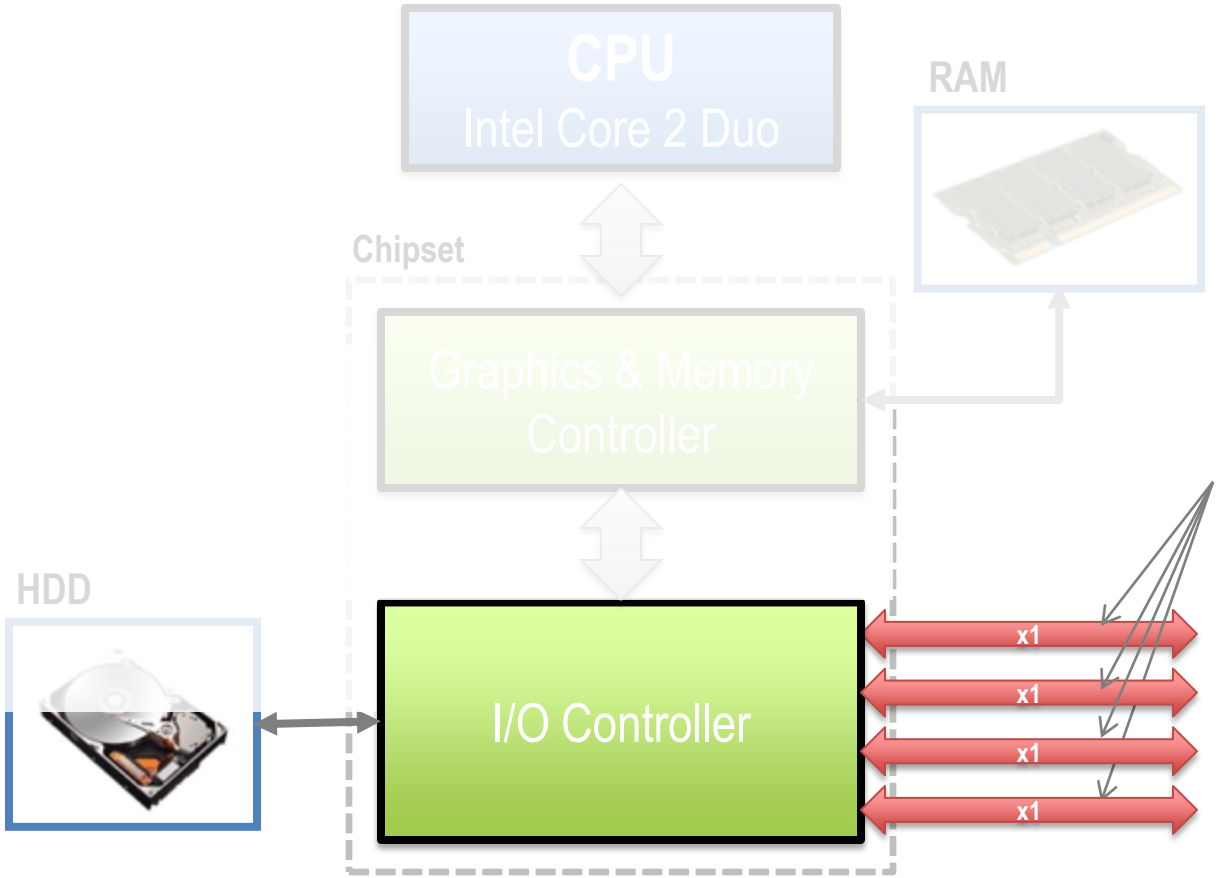

Four x1 PCI Express Links Total System Bandwidth of 1 GB/s

Forwarded to Chassis

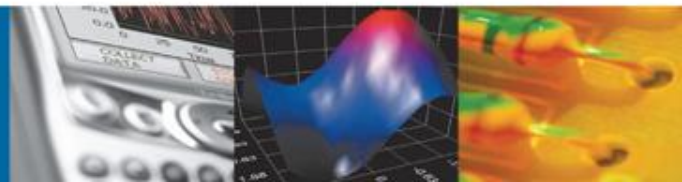

Streaming Architecture – PXI Express Chassis

NI PXIe–1062Q

Three x4 PCI Express Links
Each Capable of 1 GB/s
Throughput

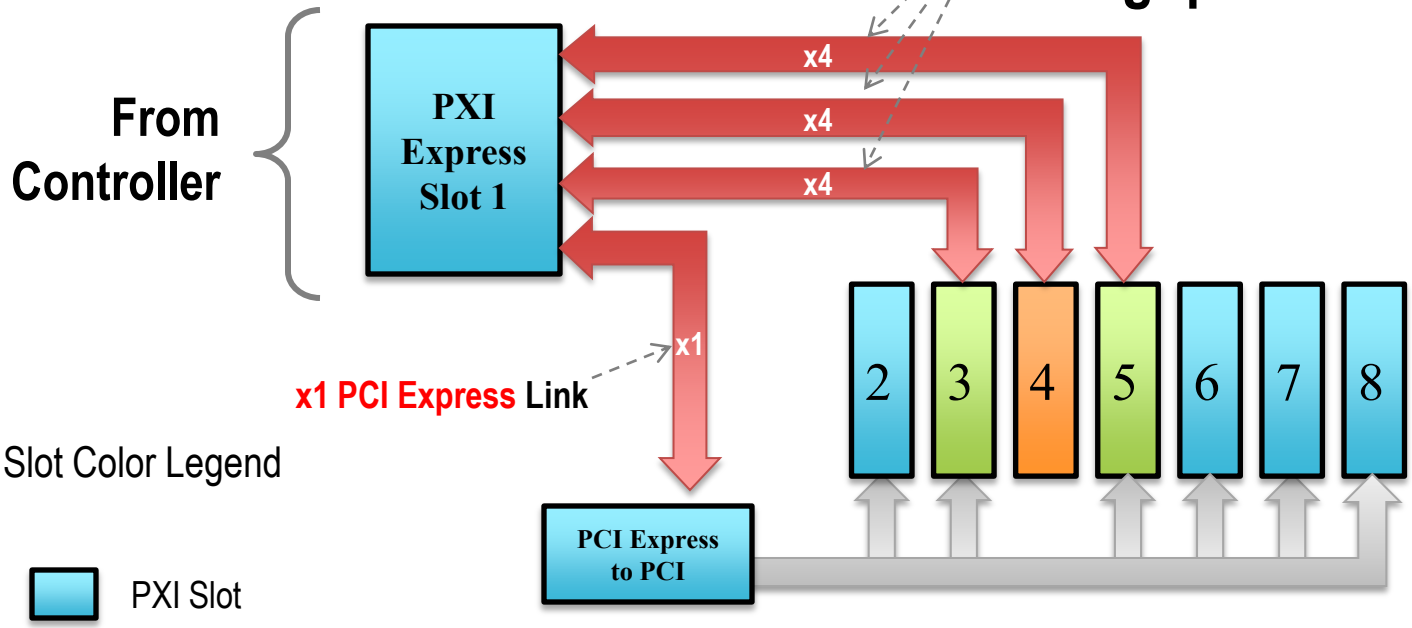

Slot Color Legend

- PXI Slot
- PXI Express/Hybrid Slot
- PXI Express System Timing Slot

NI PXIe-1075 Backplane

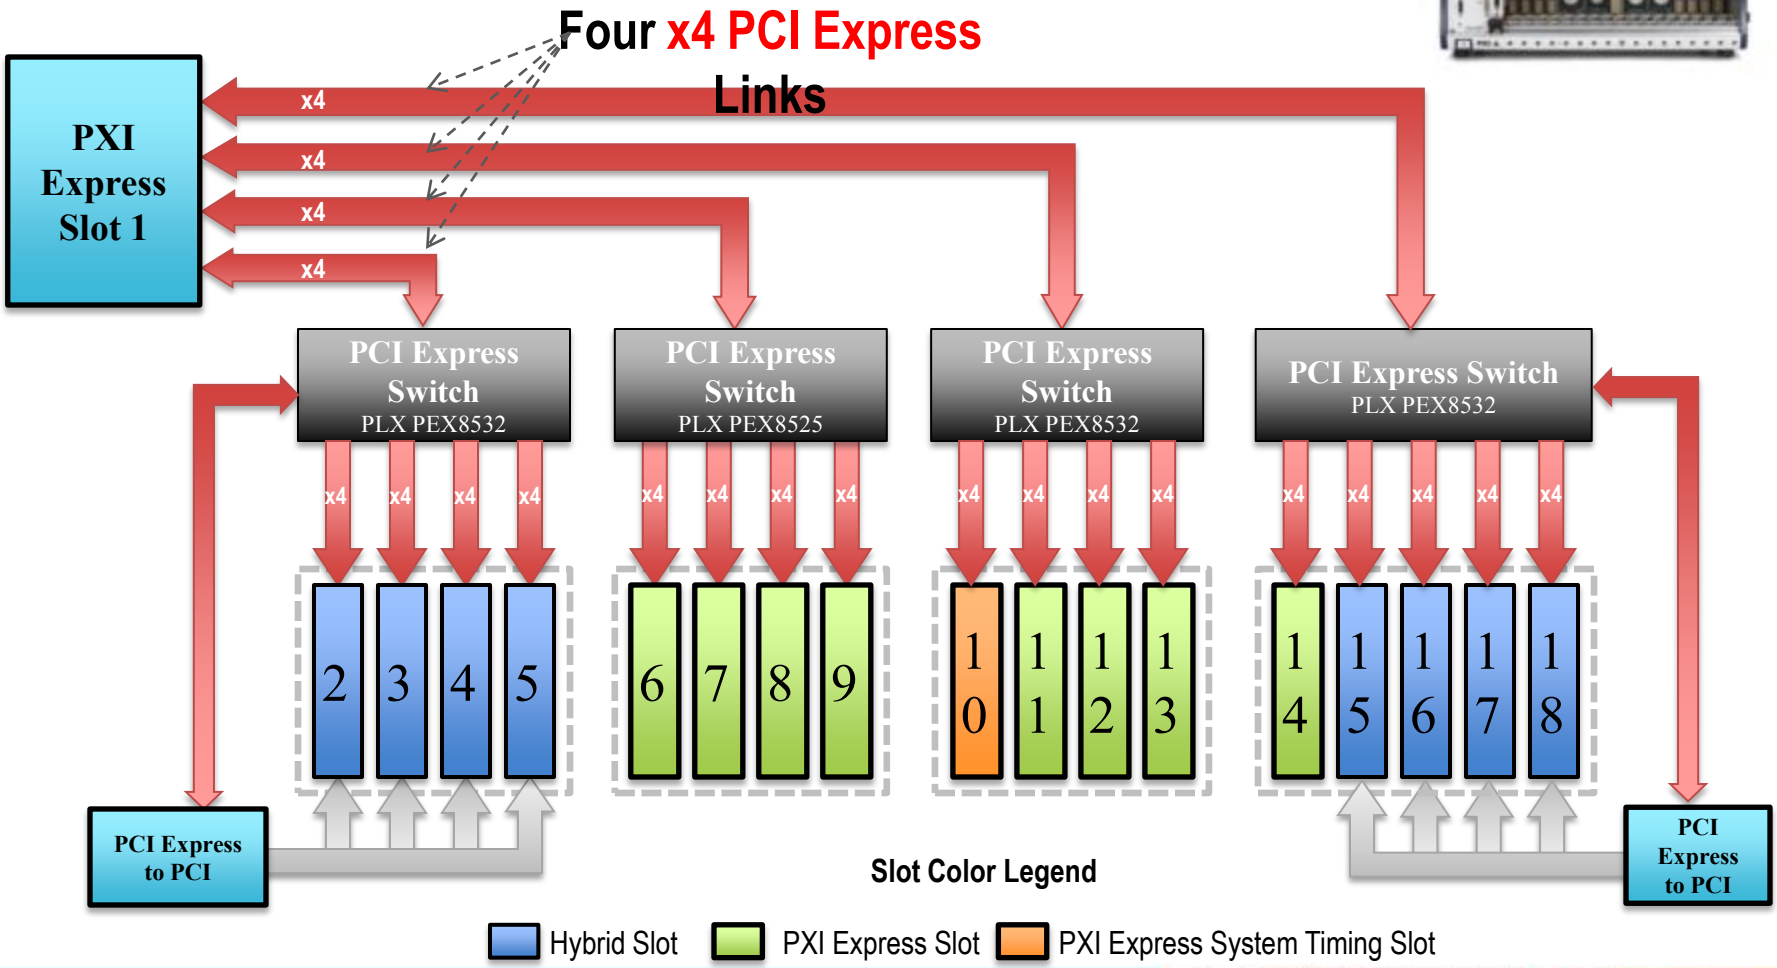

Key System Components

- Hardware Platform with High-Throughput and Low-Latency
- **High-Speed Data Storage**
  - Hard Drives (HDDs)
  - Solid-State Drives (SSDs)
- Software for Streaming to Disk at High Rates
- Streaming Front-End Instrumentation

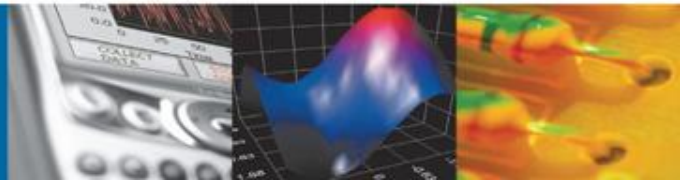

Stream To/From Disk Rates

Drive(s)	Max Rate: Write/Read (MB/s)
Laptop	30 (NI PXIe-8106 internal drive; 5,400 RPM)
IDE	57 (Western Digital 160 GB; 7,200 RPM)
SATA	62 (Western Digital 160 GB; 7,200 RPM)
SATA	75 (Seagate Barracuda 250 GB; 7,200 RPM)

- Most hard drive manufacturers do not specify streaming rates
- Specifications beyond the interface (SATA, PATA, IDE) dictate hard drive performance
  - Seek times (ms)
  - Rotational speed (RPM)
  - Buffer size (MB)
  - Density
- Benchmarking is the only guarantee

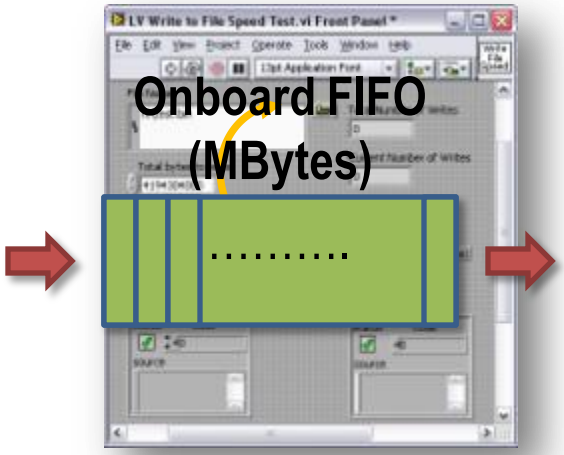

Disk Performance

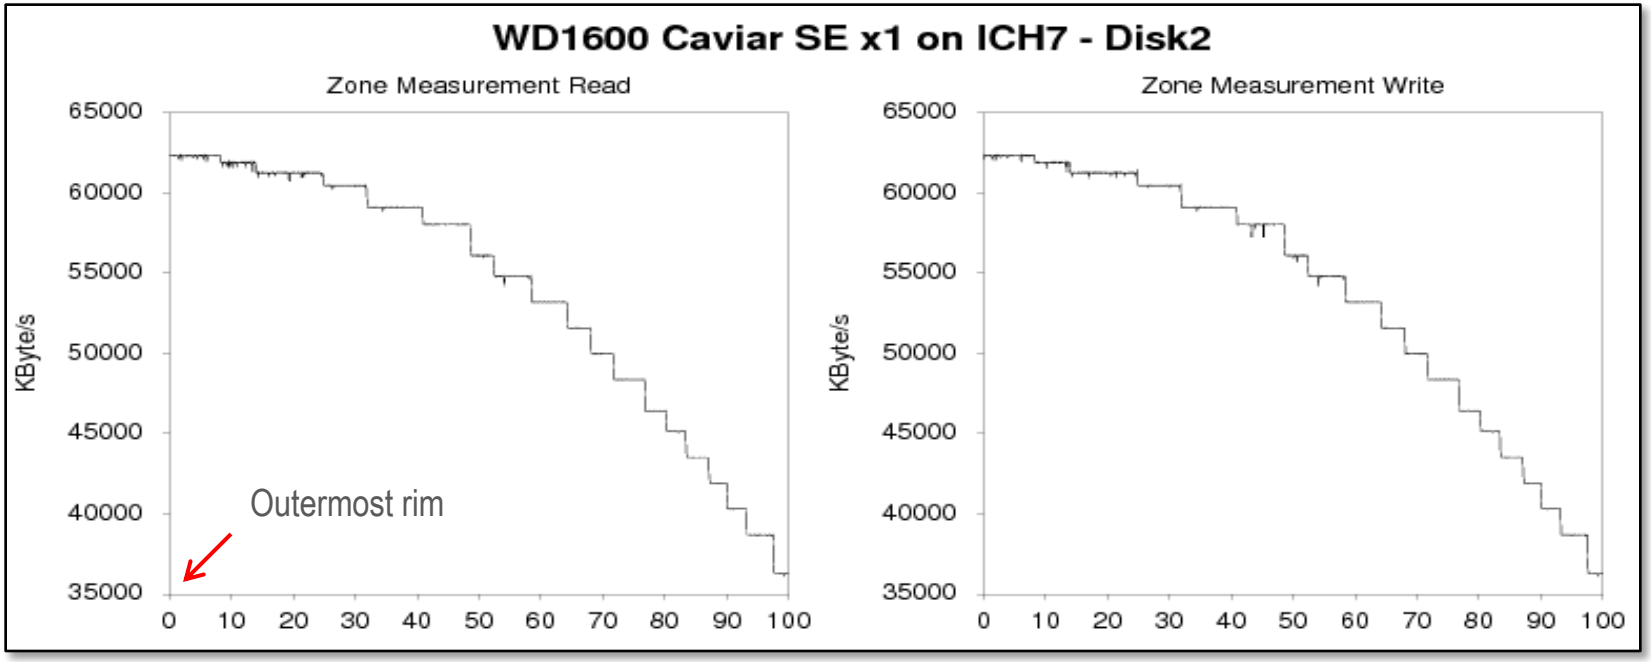

- Outer rim faster, inner rim slower
- 62 MB/s at outer rim, 36 MB/s at inner rim
- Windows OS allocates file space from outer rim inward
- True for most RAID arrays as well

Hard Drive Streaming Performance

- **Formatting**
  - Allocation unit size – the larger the better
  - Use “quick format”
- “**Write Caching**” must be turned ON
- Turn off “**System Restore**” and “**Recycle Bin**”
- File location on disk and fragmentation

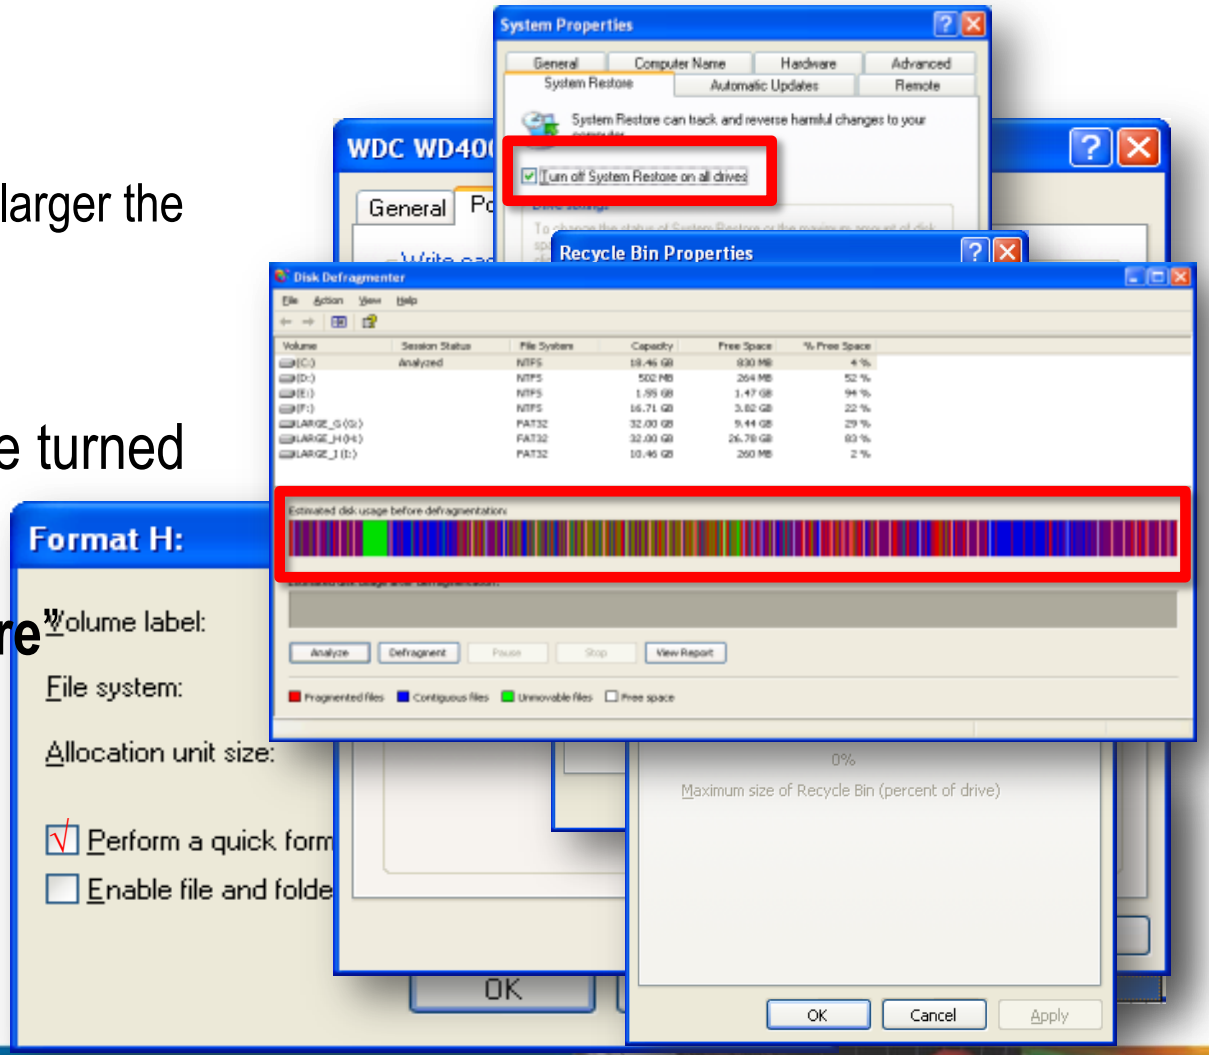

Understanding SSDs

Performance versus Capacity

- Two types of SSD Drives: **SLC** and **MLC**

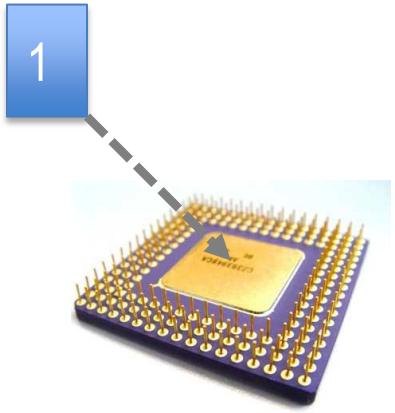

Single-Level Cell
(SLC)

Pros: Performance, Life
Cons: Low-Capacity, Cost

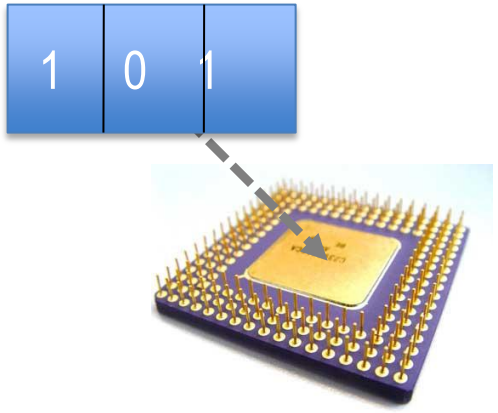

Multilevel Cell
(MLC)

Pros: Capacity, Cost
Cons: Performance, Life

Understanding SSDs

Performance versus Capacity

20 GB Write on 32 GB **SLC** SSD

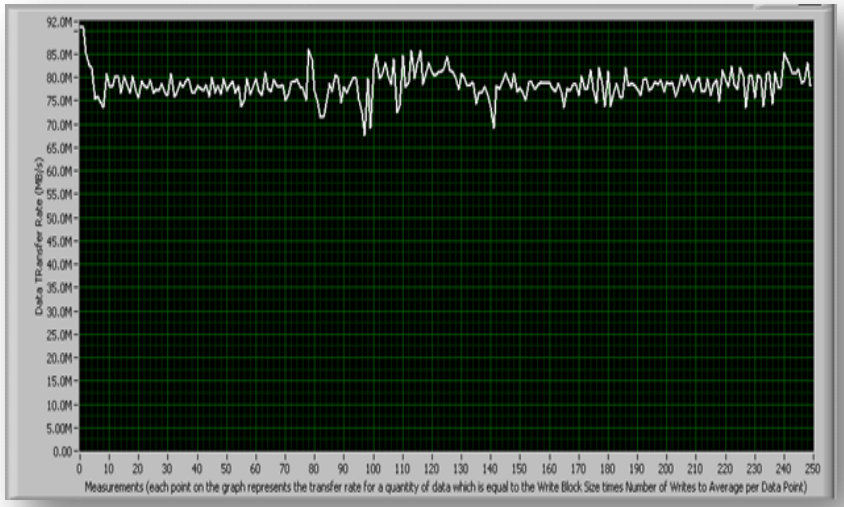

160 GB Write on 256 GB **MLC** SSD

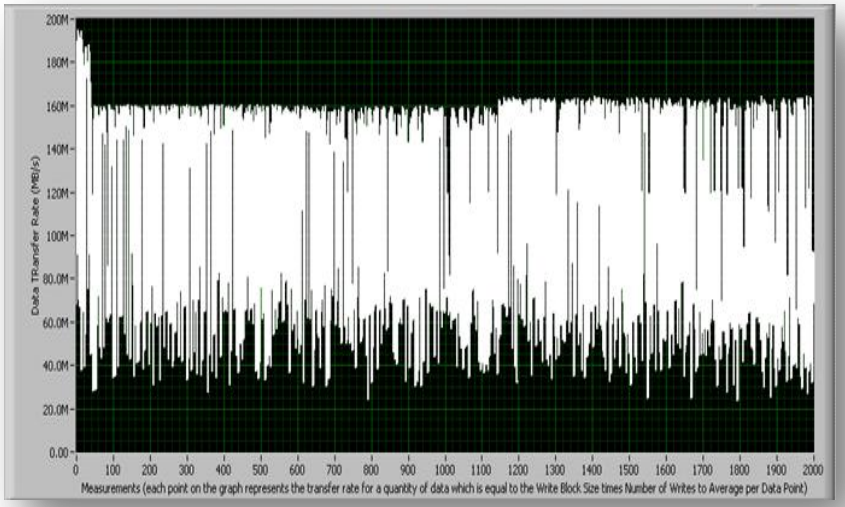

What Is RAID?

Redundant Array of Independent Drives, is a general term for mass storage schemes that split or replicate data across multiple hard drives.

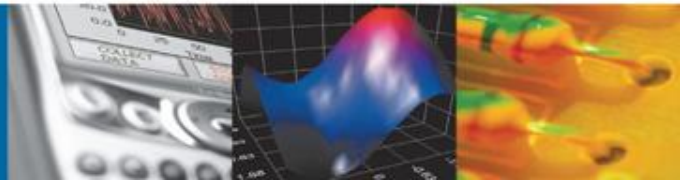

Raid 0

RAID 0 → Striping without redundancy

- Improved speed over streaming to a single hard drive
- Unimproved system reliability
- Transparently supported by Windows OS

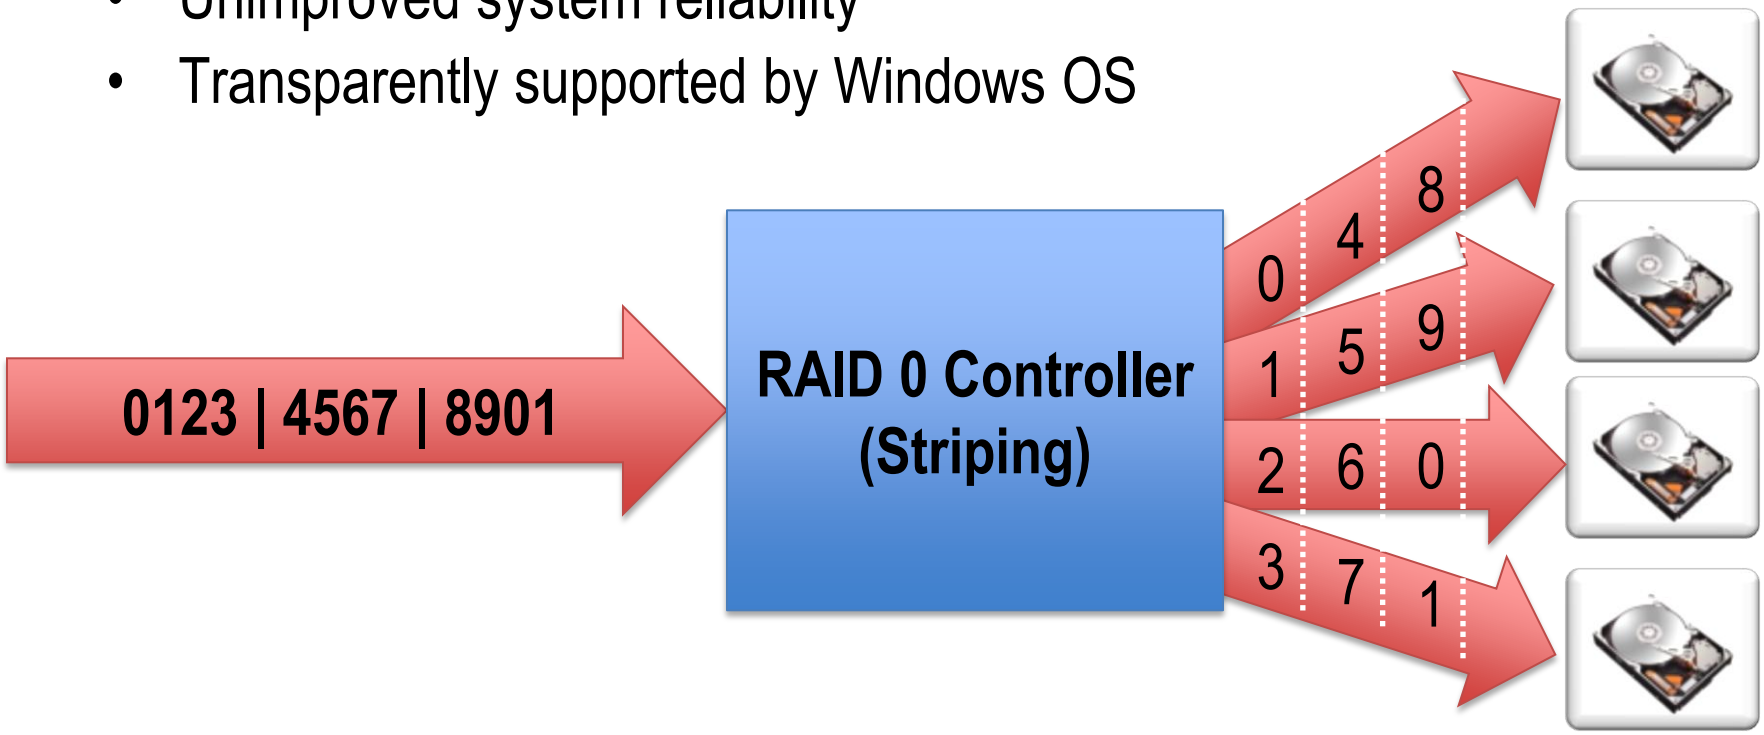

Raid 1

RAID 1 → Mirrored (redundancy)

- 100% data redundancy
- No write speed increase over single disk
- Highest overhead of all raid configurations

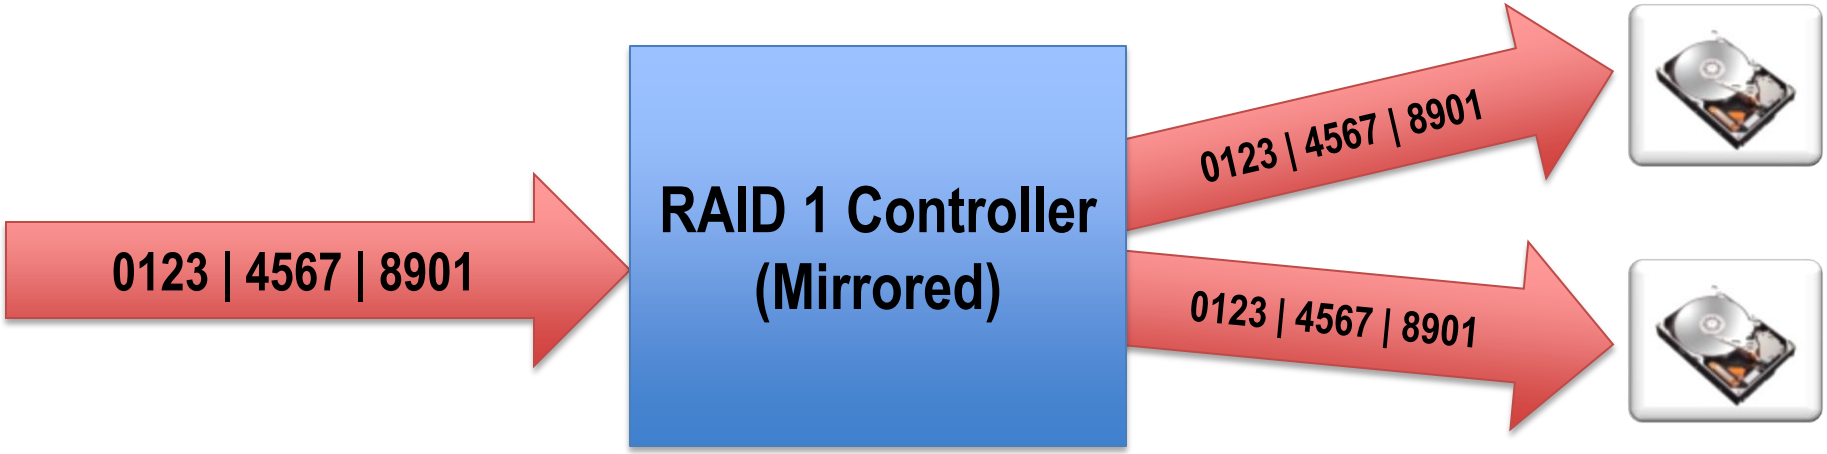

Raid 5

RAID 5→ Distributed parity

- Very efficient does not require additional disks
- Can only tolerate one drive failure
- Poor performance with small files

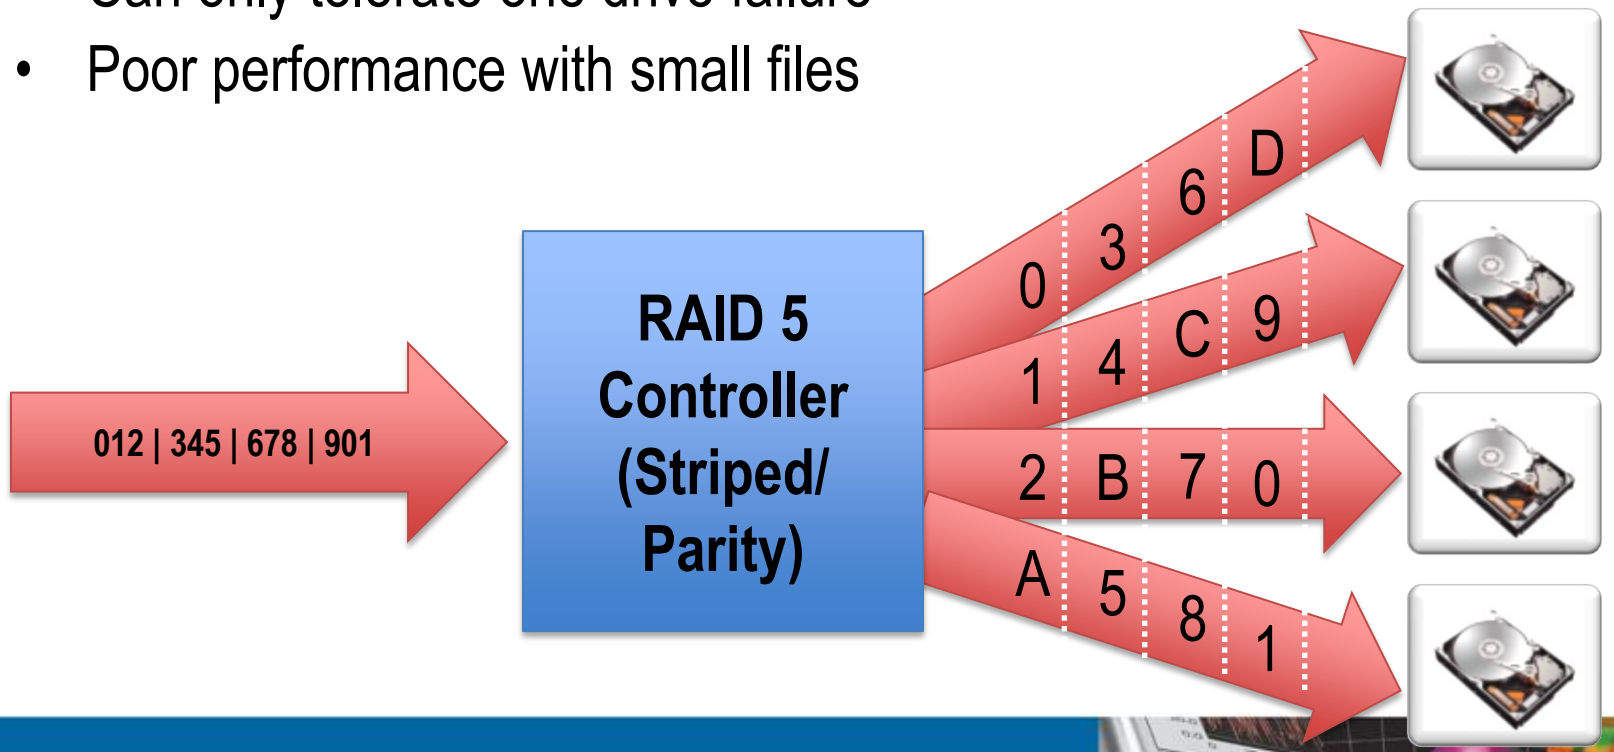

Raid 1+0

RAID 1+0 or 10 → Striping and mirroring

- Highest performance with data redundancy
- Can sustain multiple drive failures
- Configuration requires twice the hard drives

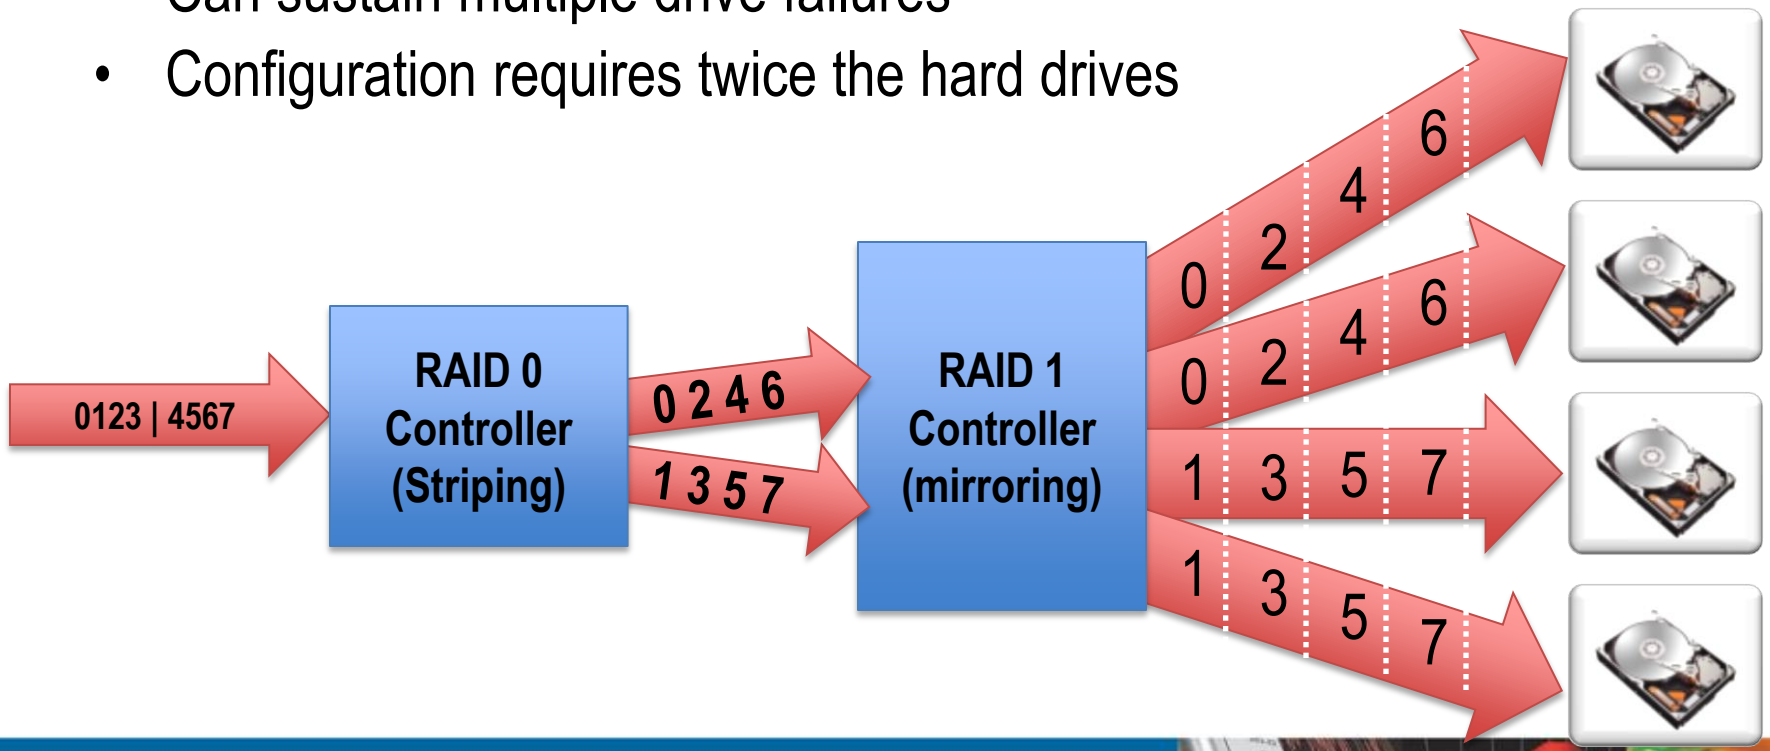

Stream To/From Disk Rates

Drive(s)	Write/Read (MB/s)	Rate Types
Laptop	30 (NI PXIe-8103 internal drive; 5,400 RPM)	Peak
IDE	57 (Western Digital 160 GB; 7,200 RPM)	Peak
SATA	62 (Western Digital 160 GB; 7,200 RPM)	Peak
SATA	75 (Seagate Barracuda 7,200.10; 250 GB)	Peak
2 RAID	114/127 (NI PXI-8351 1U Rack Mount Controller)	Peak
4 RAID	200+/200+ (NI HDD-8263, NI 8353, NI 8260)	Sustained
12 RAID	600/600 (NI 8264 RAID Controller)	Sustained

Outer rim rates. Cannot be sustained across the whole drive.

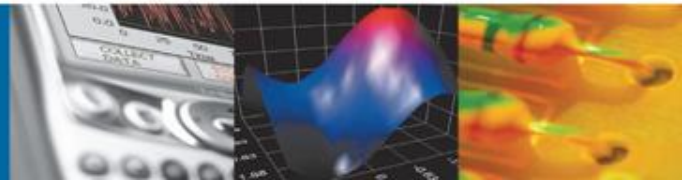

Data Streaming Products

Rackmount (External)

NI HDD-8263

- 200 MB/s
- 1 TB
- More than 1 hour at 100 MS/s

NI HDD-8264

- 600 MB/s
- 3 TB
- More than 3 hours at 100 MS/s

In-Chassis (Internal)

NEW!

NI HDD-8260

- 200 MB/s
- 1 TB
- More than 1.5 hour at 100 MS/s

- 3-slot wide, 4-drive
- SSD option (128 GB) available
- Software RAID

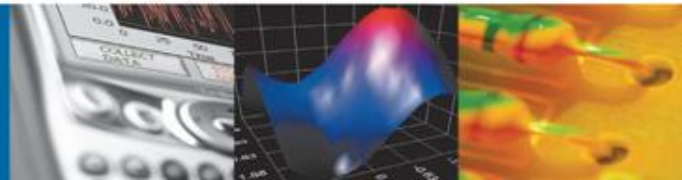

NI 8260

In-Chassis High-Speed Storage Module

- For PXI Express systems
- 3-slot wide storage module
- 4-drive software RAID
- 200 MB/s
- HDD version: 1 TB
- SSD version: 128 GBs

Key System Components

- Hardware Platform with High-Throughput and Low-Latency
- High-Speed Data Storage
  - Hard Drives (HDDs)
  - Solid-State Drives (SSDs)
- **Software for Streaming to Disk at High Rates**
- Streaming Front-End Instrumentation

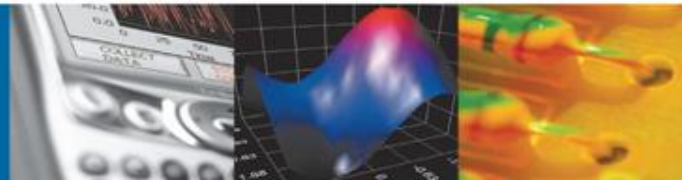

Using Data Streaming Products

- The RAID hard drives appear as logical partition in Windows OS
- LabVIEW 8.5.1 or later

  Win32 file I/O VIs
- LabVIEW 8.6 and later

  Built-in file VIs
- LabVIEW 2009

  Can Use LabVIEW 2009 TDMS

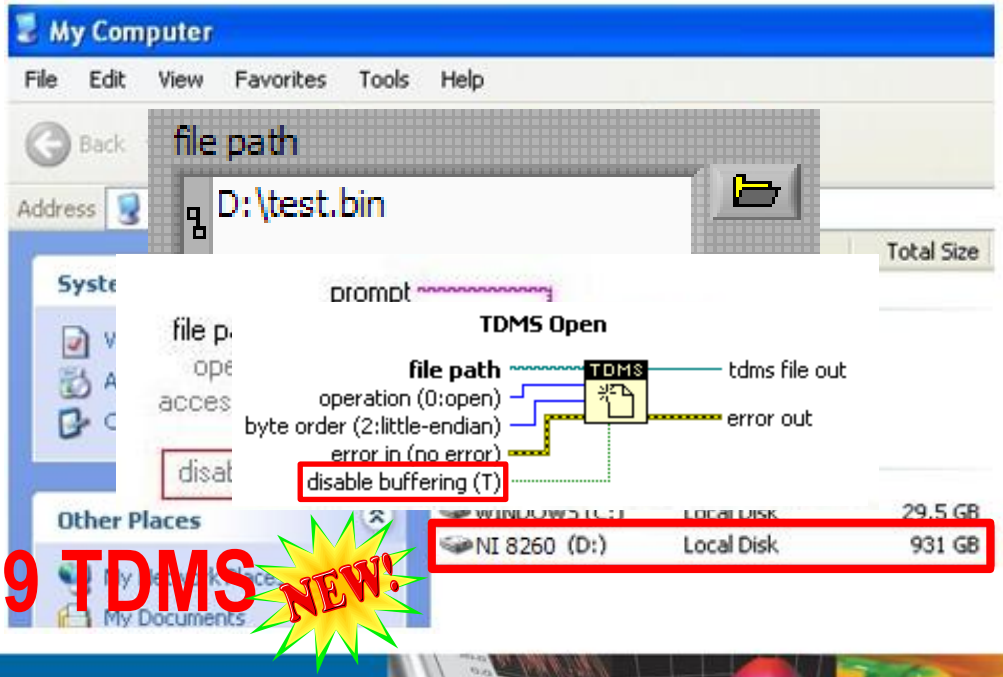

LabVIEW Programming Structure

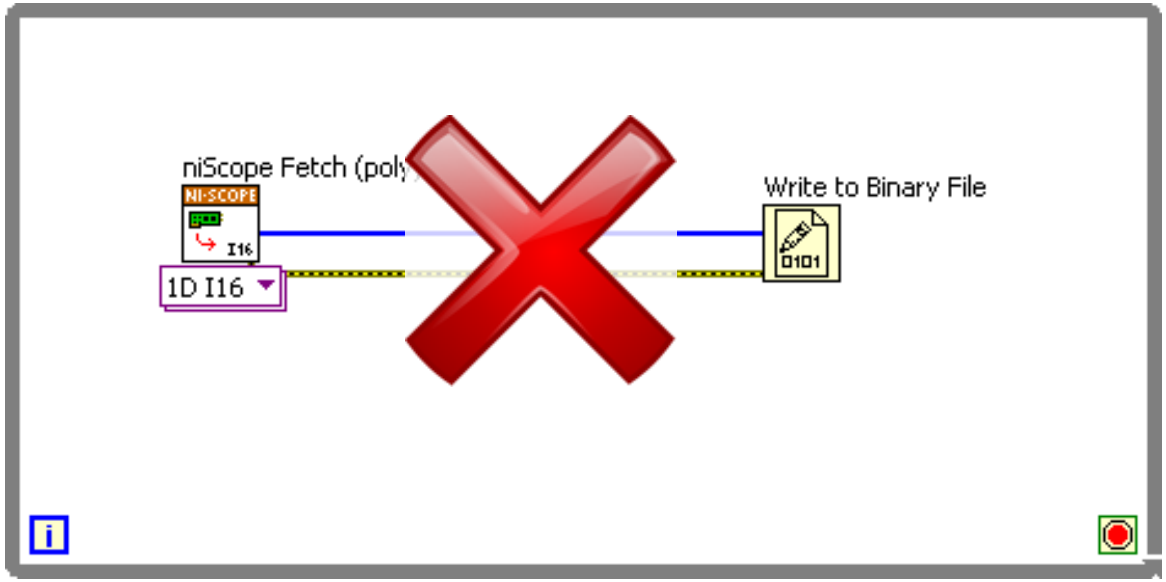

T_1 T_2

Loop Time = $T_1 + T_2$

Use Multithreading

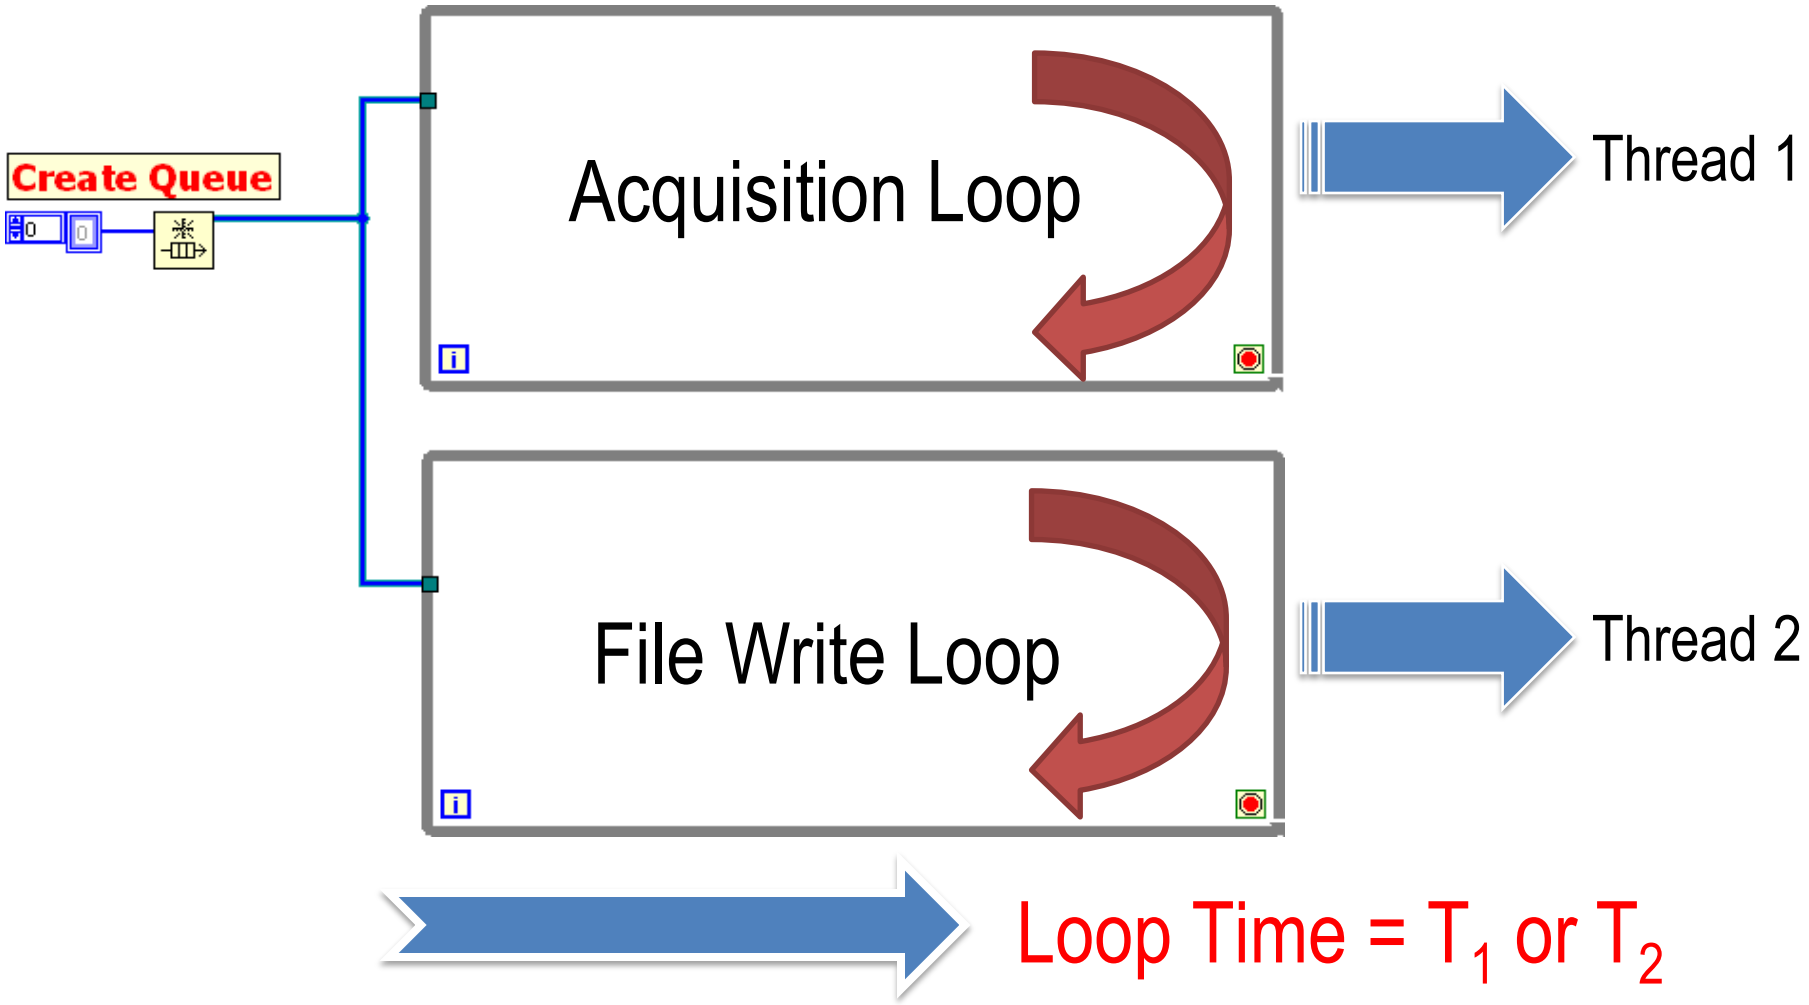

Producer Consumer Loops

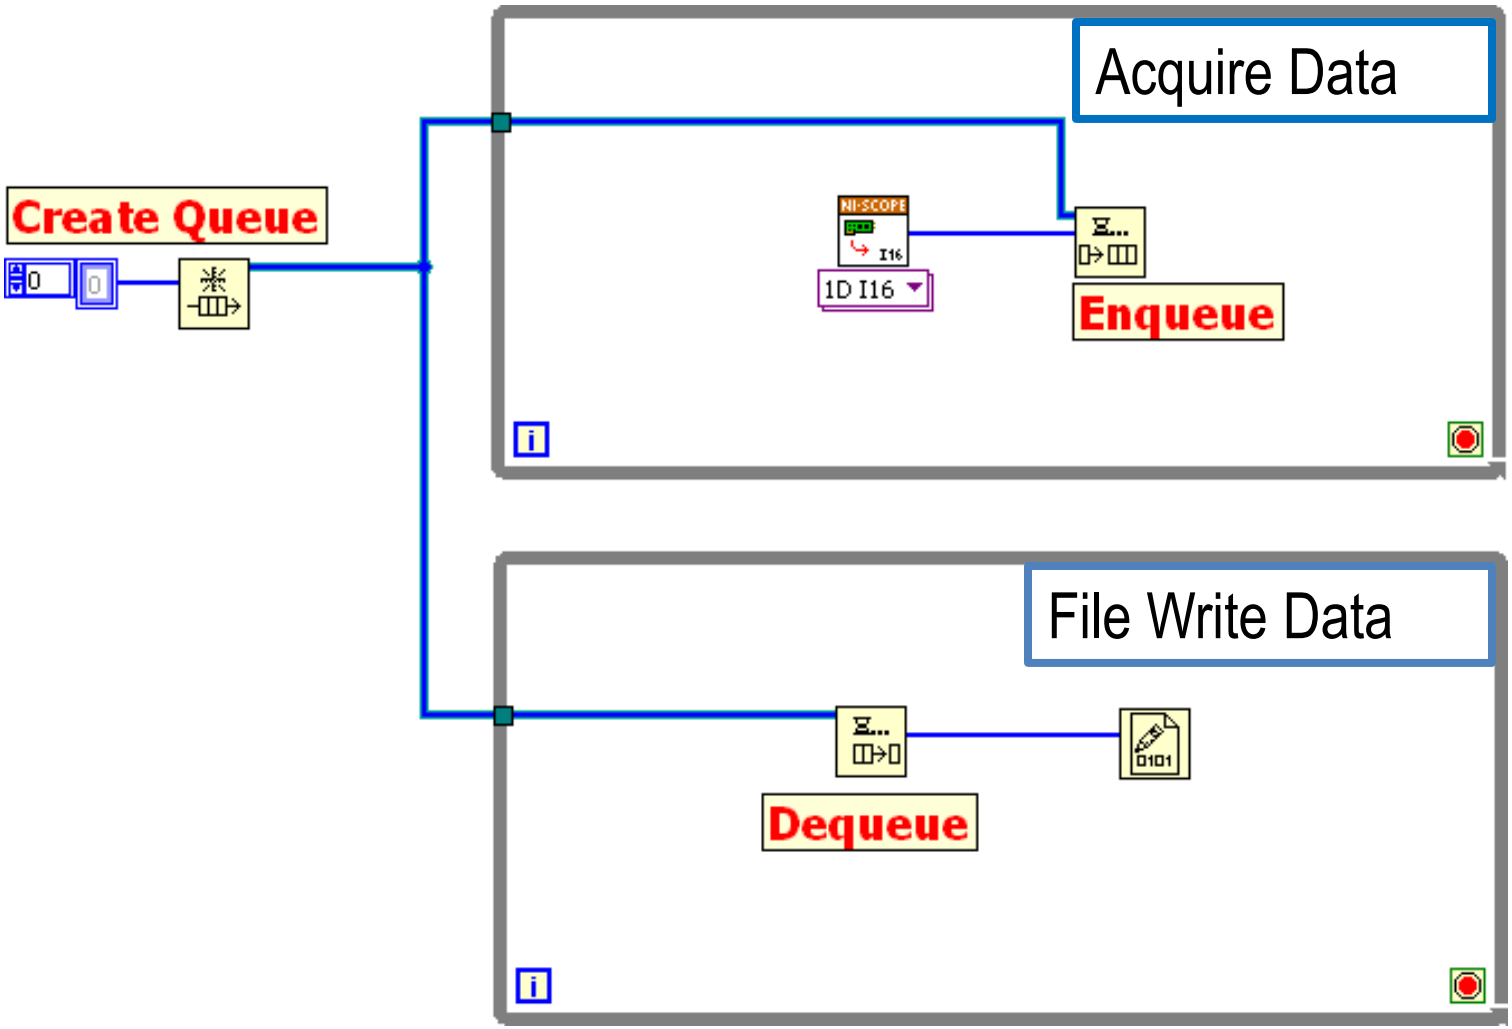

Data Types

- Reduce file size
  - 1 I16 sample = 16 bits = 2 bytes
  - 1 DBL sample = **64 bits** = **8 bytes = 4X increase in bandwidth**

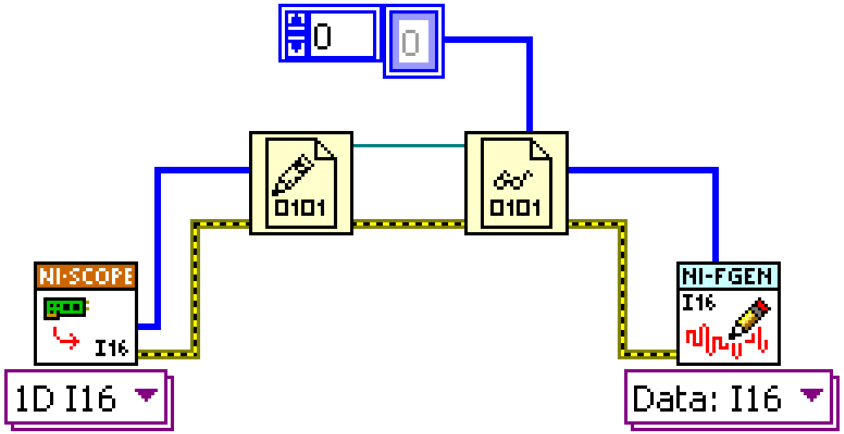

Demo 1

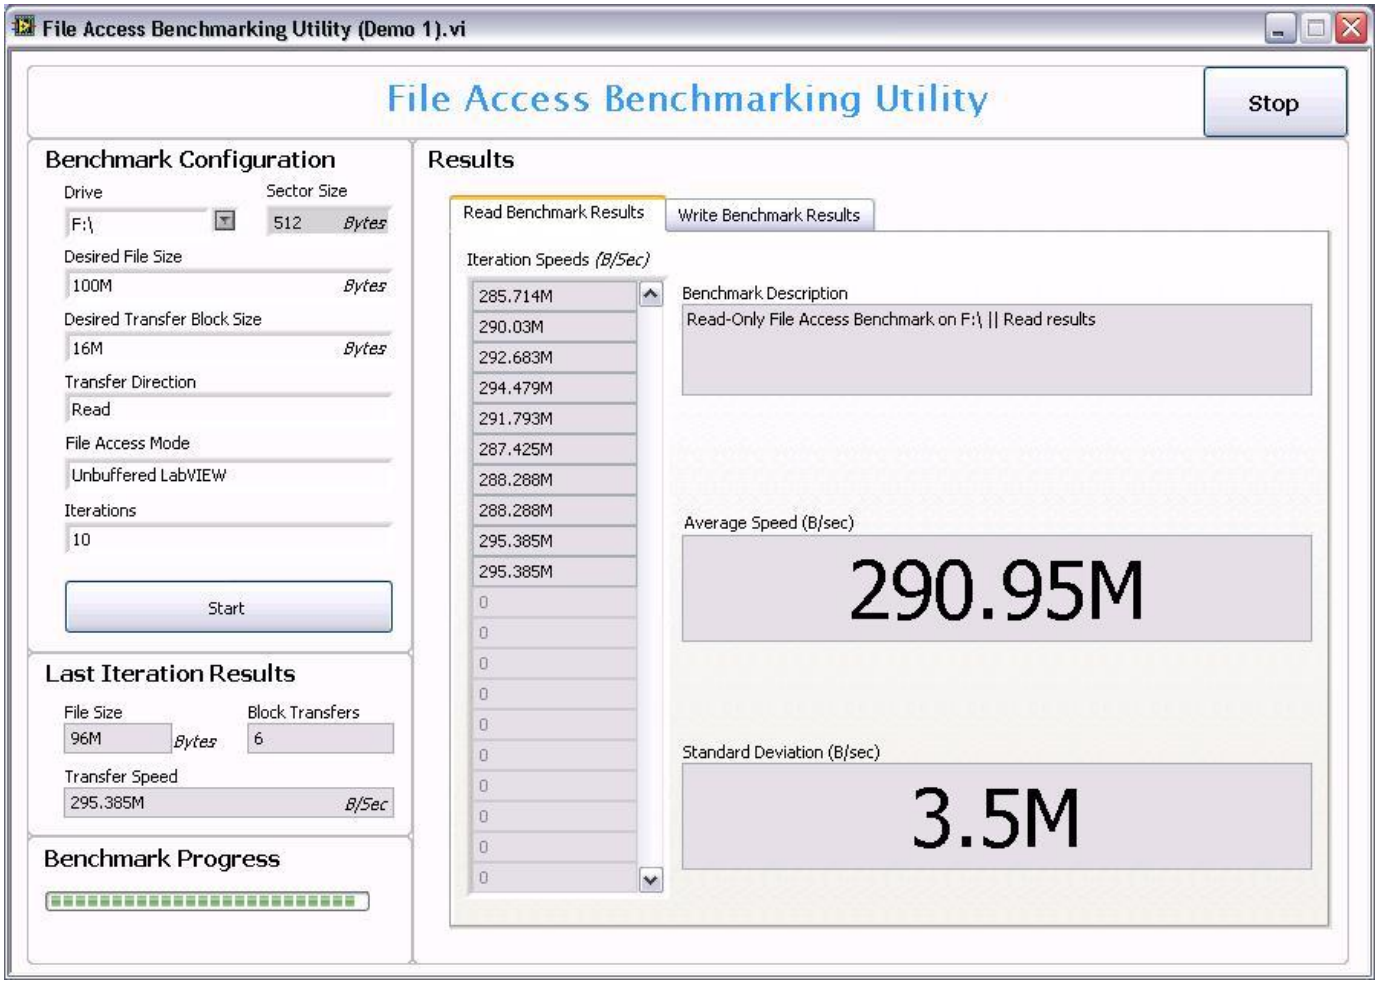
File Access Benchmarking Utility (Demo 1).vi
File Access Benchmarking Utility
Stop
Benchmark Configuration
Drive
F:\
Sector Size
512 Bytes
Desired File Size
100M Bytes
Desired Transfer Block Size
16M Bytes
Transfer Direction
Read
File Access Mode
Unbuffered LabVIEW
Iterations
10
Start
Last Iteration Results
File Size
96M Bytes
Block Transfers
6
Transfer Speed
295.385M B/Sec
Benchmark Progress
Results
Read Benchmark Results
Write Benchmark Results
Iteration Speeds (B/Sec)
285.714M
290.03M
292.683M
294.479M
291.793M
287.425M
288.288M
288.288M
295.385M
295.385M
0
Benchmark Description
Read-Only File Access Benchmark on F:\ || Read results
Average Speed (B/sec)
290.95M
Standard Deviation (B/sec)
3.5M

Key System Components

- Hardware Platform with High-Throughput and Low-Latency
- High-Speed Data Storage
  - Hard Drives (HDDs)
  - Solid-State Drives (SSDs)
- Software for Streaming to Disk at High Rates
- **Streaming Front-End Instrumentation**

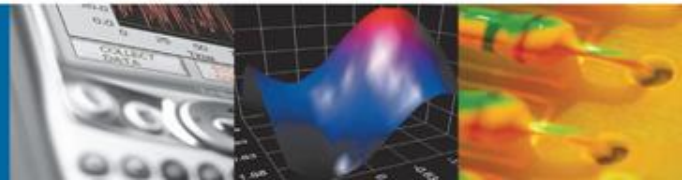

Streaming Modules

PXIe-6544/45 DIO
- 100/200 MHz
- 32 lines
- 1.2,1.51.8,2.5,3.3 V
- 660 MB/s

PXIe-5442 Arb
- 100 MS/s
- 43 MHz, 16-bit
- 40 MHz DUC
- 200 MB/s

PXIe-5122 Digitizer
- 2 channel
- 100 MS/s
- 100 MHz, 14-bit
- 400 MB/s

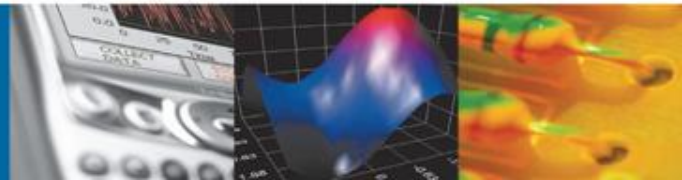

Streaming Modules

PXIe-5673 VSG

- 85 MHz to 6.6 GHz
- 100 MHz Bandwidth
- 125 MS/s IQ Rate
- 500 MB/s

PXIe-5663 VSA

- 10 MHz to 6.6 GHz
- 50 MHz Bandwidth
- 62.5 MS/s IQ Rate
- 250 MB/s

Streaming Modules

PXIe-5450 IQ Generator

- 400 MS/s, 16-bit, dual-channel
- 145 MHz
- 600 MB/s dual-channel
- 360 MB/s single-channel

PXIe-5622 IF Digitizer

- 150 MS/s, 16-bit
- 3-250 MHz
- 60 MHz DDC
- 300 MB/s

Streaming Modules

PXIe-8234 GigE Interface

- Dual-port gigabit Ethernet
- NI Vision software
- 250 MB/s

Streaming Performance

- Most operations are possible with a direct link to the PXI Express controller (no switches)
- Chassis and controller set maximum system bandwidth
- Module location and type (input or output) are critical factors

Instrument Behavior: Input

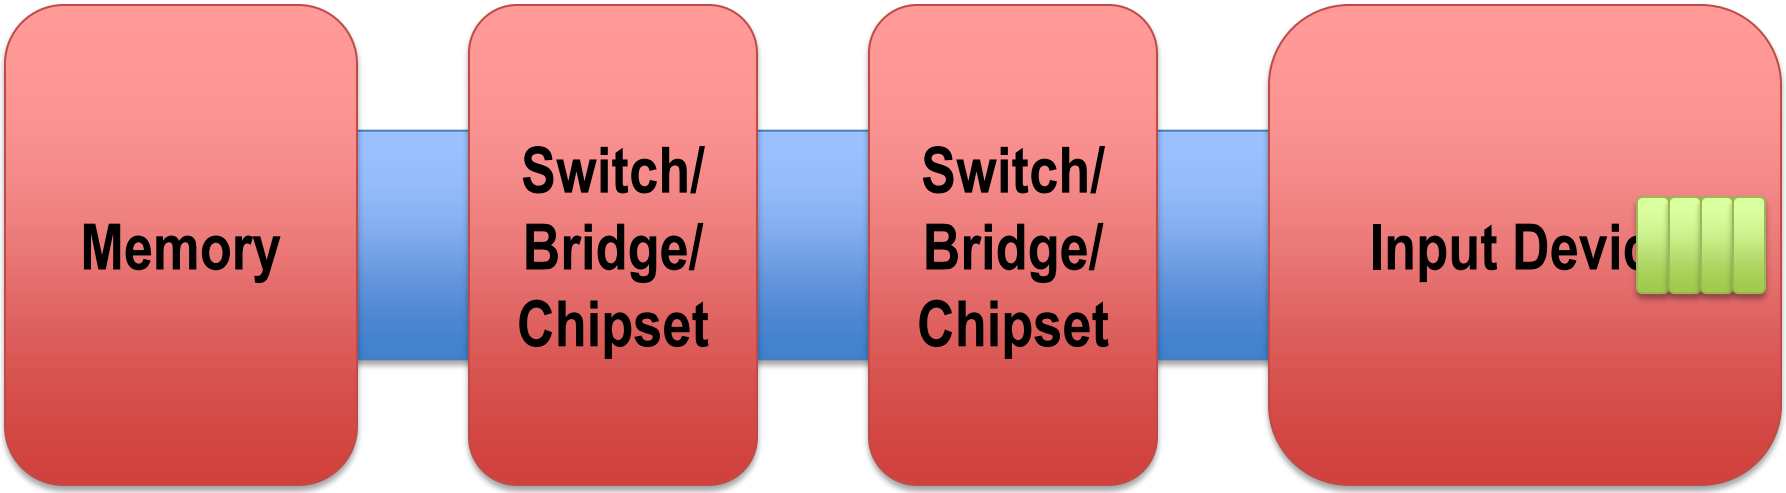

Data

Instrument Behavior: Output

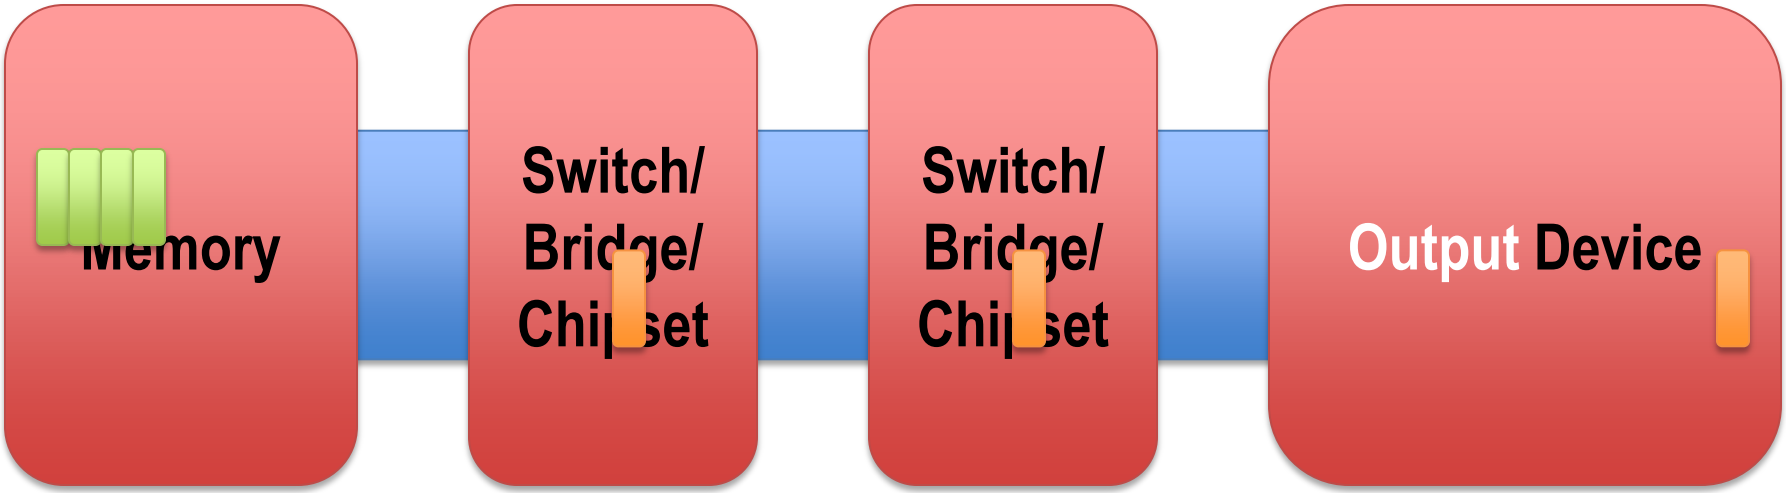

Read Request

Data

High-Level Guidelines

- Streaming devices should be given their own dedicated link or switch, if possible
- RAID arrays should be the only streaming device on a dedicated controller link, whether or not they are behind a switch
- Modules with opposite streaming direction (input and output) should ***not*** be grouped behind the same switch
- Devices based on PXI Express NI-DAQmx (small onboard memory) should share their own switch

Lower than 400 MB/s aggregate bandwidth (input + output);
these should not be a problem

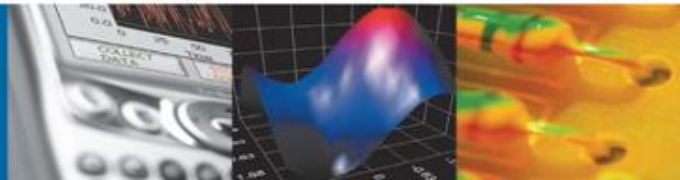

A High-Performance Application

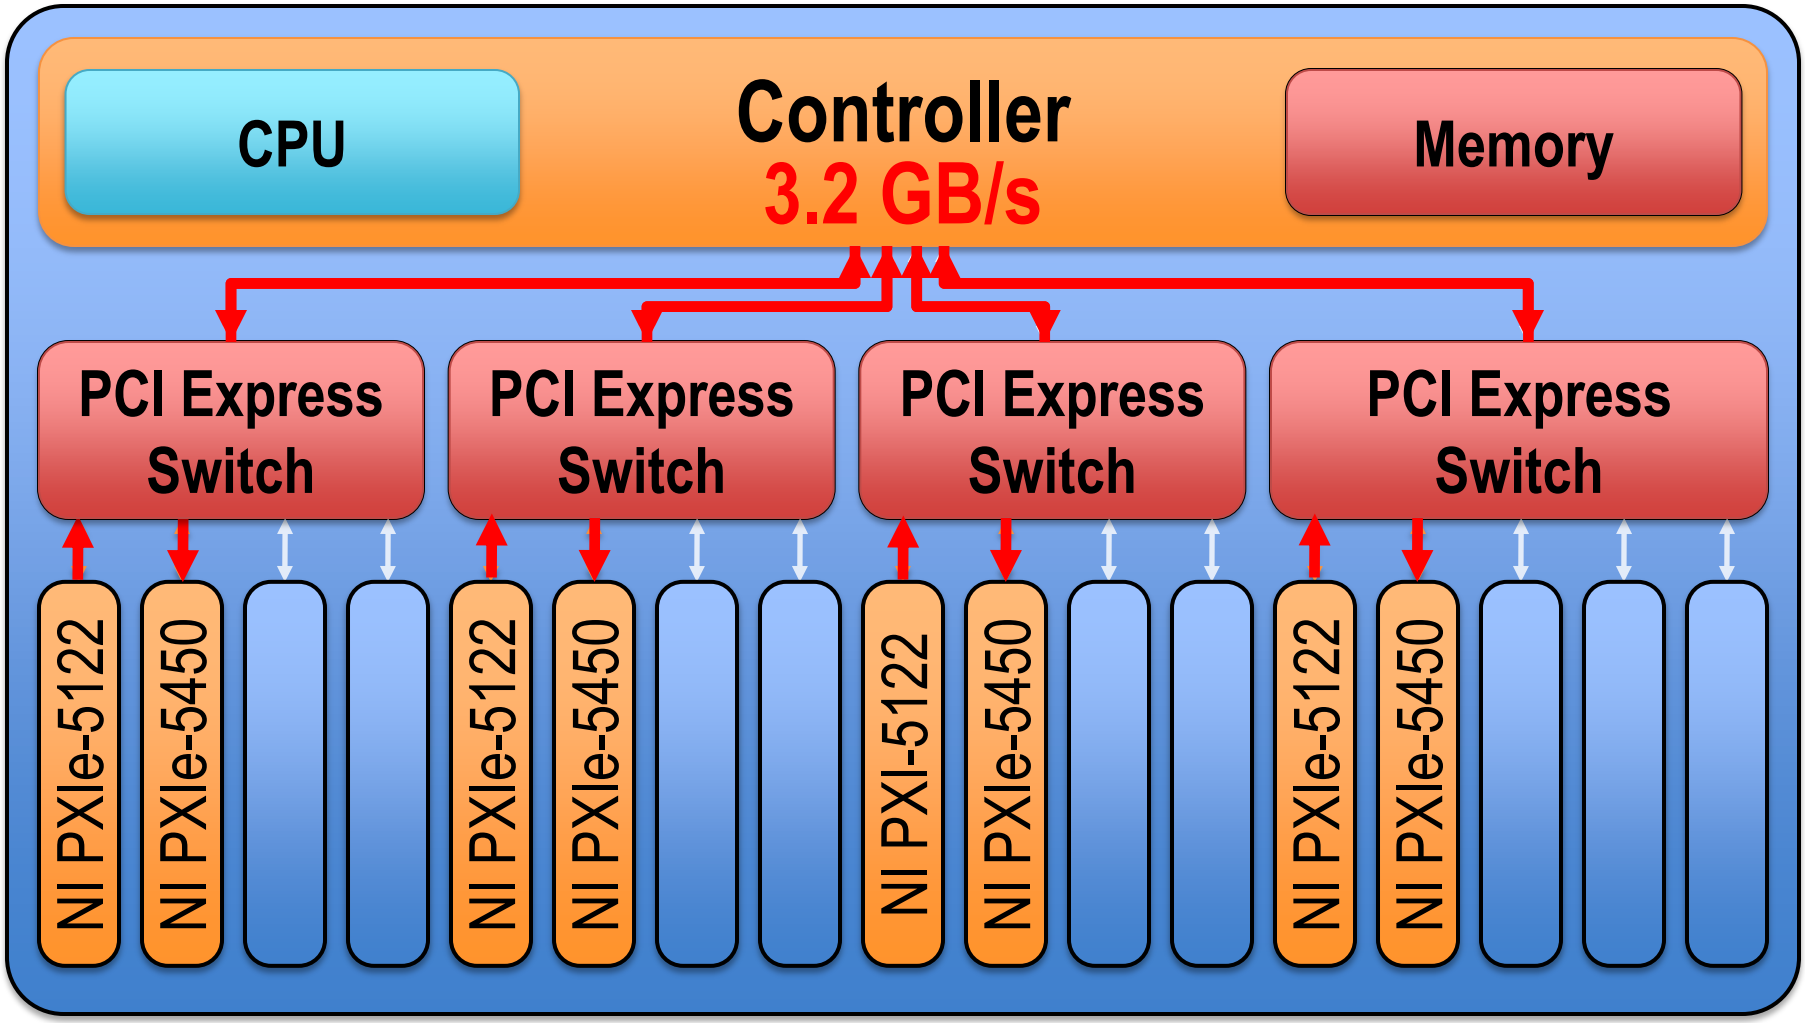

Demo 2

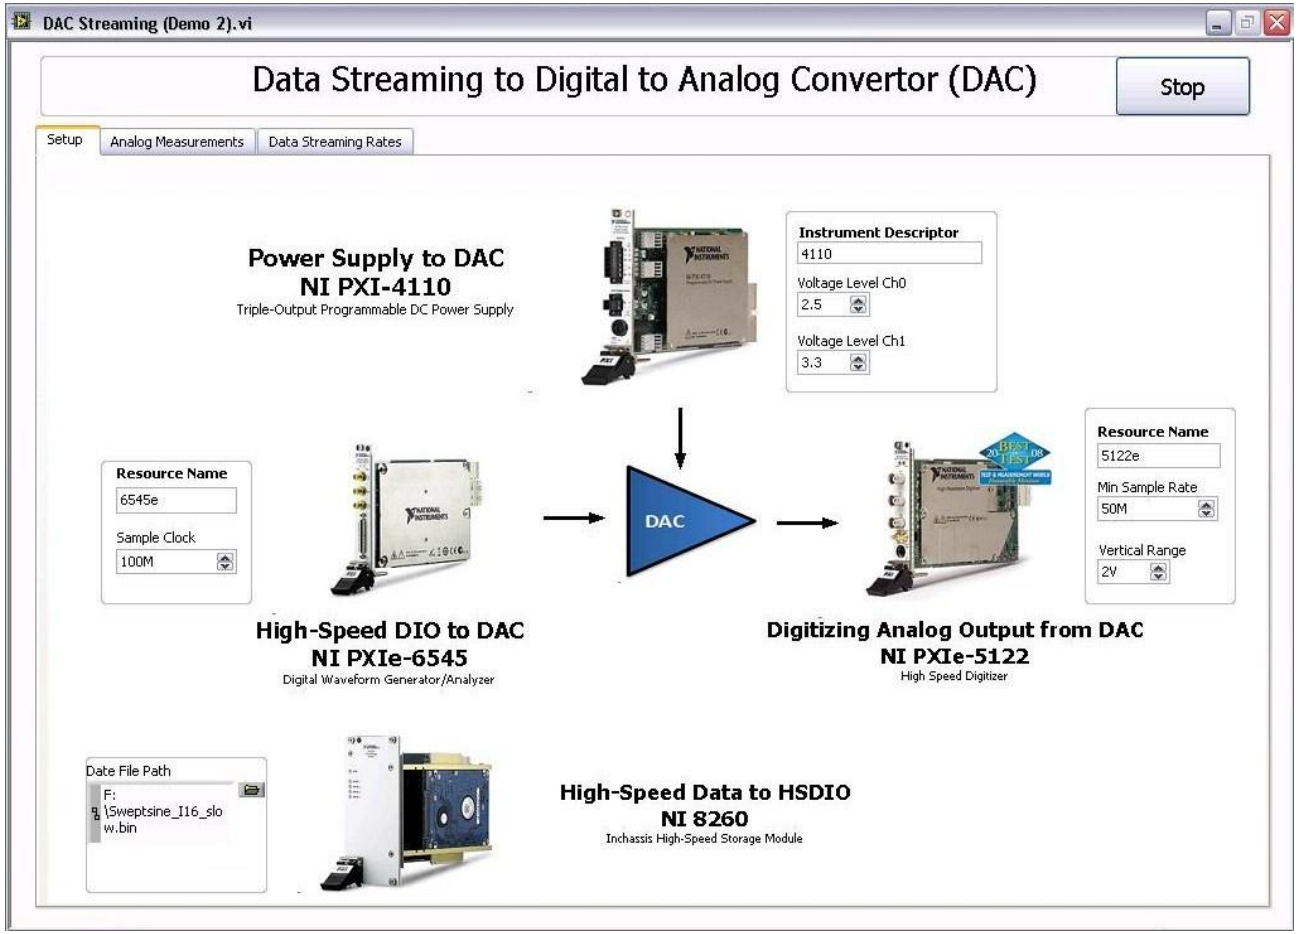
DAC Streaming (Demo 2).vi
Data Streaming to Digital to Analog Convertor (DAC)
Stop
Setup
Analog Measurements
Data Streaming Rates
Power Supply to DAC
NI PXI-4110
Triple-Output Programmable DC Power Supply
Instrument Descriptor
4110
Voltage Level Ch0
2.5
Voltage Level Ch1
3.3
Resource Name
6545e
Sample Clock
100M
DAC
Resource Name
5122e
Min Sample Rate
50M
Vertical Range
2V
High-Speed DIO to DAC
NI PXIe-6545
Digital Waveform Generator/Analyzer
Digitizing Analog Output from DAC
NI PXIe-5122
High Speed Digitizer
Date File Path
F:
\Sweptsine_I16_slo
w.bin
High-Speed Data to HSDIO
NI 8260
Inchassis High-Speed Storage Module

Questions?