



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2016

# Enterprise Storage RAS augmented by native Intel Platform Storage Extensions (PSE)

**Tanveer Alam**  
**Intel Corporation**

# Objective

A successfully managed, RAS (Reliability, Availability, Service-ability) capable enterprise and cloud storage system, are the keys to building a highly predictable and resilient storage strategy for any enterprise or mid-sized business. However, to implement this strategy, it takes the right hardware and software building blocks. Intel<sup>®</sup> Xeon<sup>®</sup> Processor based processors comprise of some of the key fundamental hardware technology extensions that enables enhanced storage system resiliency when used in combination with other higher domain hardware or software such as RAID6 and other compatible scalable redundancy solutions. Two of the most salient Intel Platform Storage Extensions (PSE) include PCIe based Intel Non-Transparent Bridge (NTB) and Intel Asynchronous DRAM Refresh (ADR), Intel<sup>®</sup> QuickData technology natively built into the Intel Xeon processors.

# Agenda

- ❑ RAS Practical Definitions and Overview
- ❑ RAS Enabling Framework
- ❑ RAS Features for Storage
- ❑ Intel Platform Storage Extension (PSE) technologies augmenting system RAS:
  - ❑ Non-transparent Bridge (NTB)
  - ❑ Asynchronous DRAM Refresh (ADR)
  - ❑ Intel® QuickData Technology (chipset DMA acceleration)
- ❑ RAS System Visibility
- ❑ Final Thoughts

# What is RAS? – Practical definitions

**RAS\*\* = Reliability, Availability, Service-ability :**  
**99.999% UPTIME (Five Nines Reliability)**

- ❑ **Reliability = Engineering/Architectural Decision**
  - ❑ **Error Detection and Self-Healing**
  - ❑ **Minimizes outage**
  - ❑ **Correct certain bits (therefore data) at all times**
- ❑ **Availability = Keep running despite problems**
  - ❑ **Reduce frequency and duration of outages**
  - ❑ **Self-Diagnosing, Self-healing (auto-workaround problems)**
  - ❑ **24/7 operation – business never stops or slows down**
- ❑ **Serviceability = Minimized outage/downtime**
  - ❑ **Avoid redundant failures and provide accurate failure diagnostics**
  - ❑ **Concurrent repair (heal) of higher failure items**
  - ❑ **Easy repair/replace/upgrade.**

# RAS: One Size Doesn't Fit All



## Cloud

Need Fault Handling

Check pointing is not used

Apps can tolerate single system failure

Fault mgmt. for improving TCO

Ex: Automated techniques to identify failed DIMM

## HPC

Needs fault handling

Check pointing heavily used

Apps can not tolerate single system failure

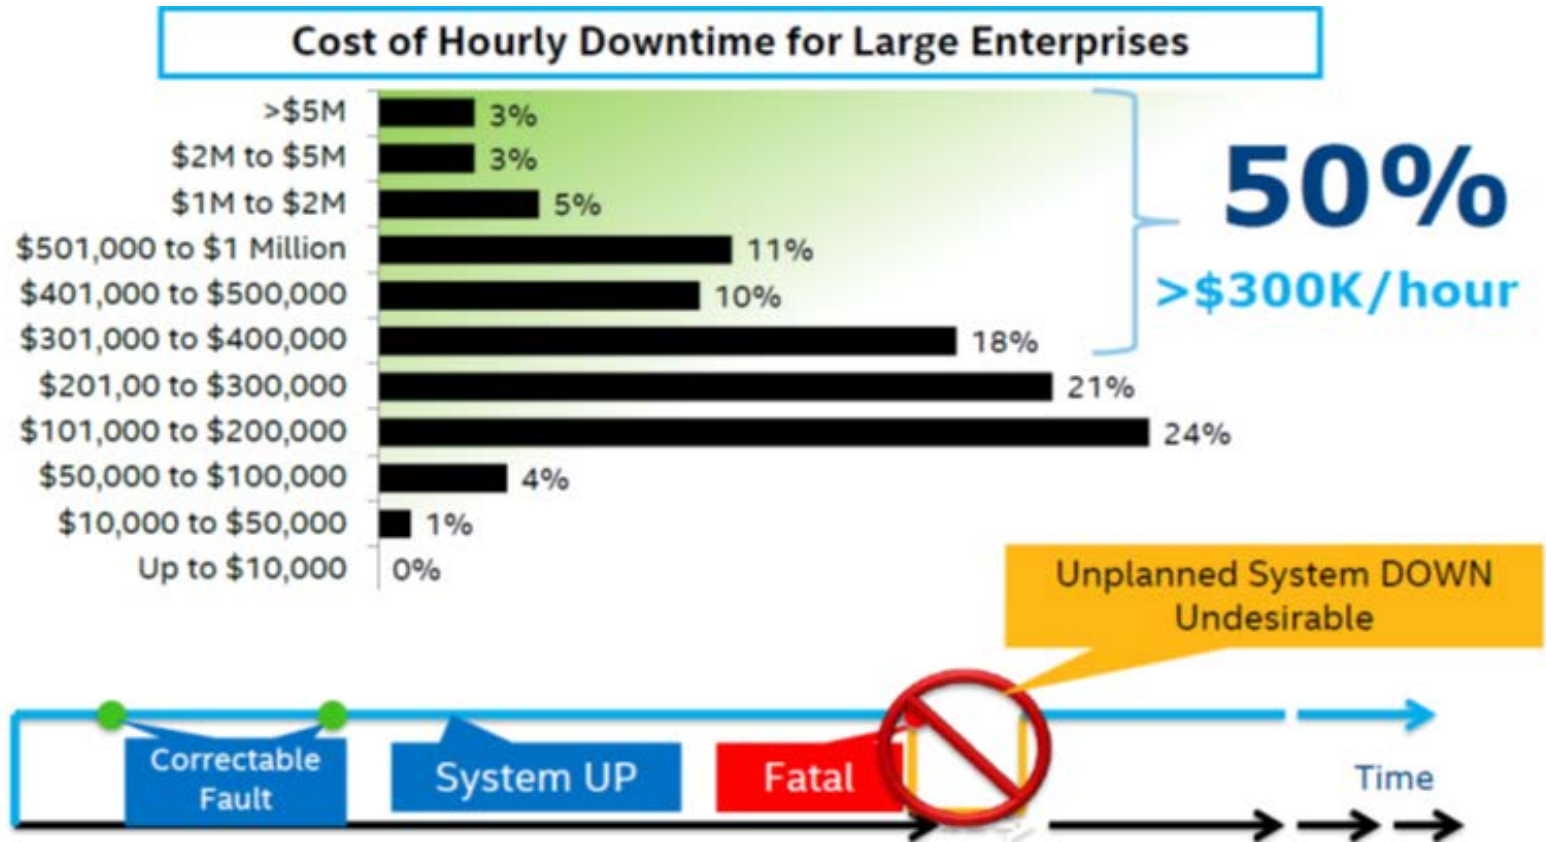
Fault mgmt. for improving Uptime

Ex: HW/FW based self-healing

**Different applications may require different RAS features**

# RAS Needed in Mission Critical Systems

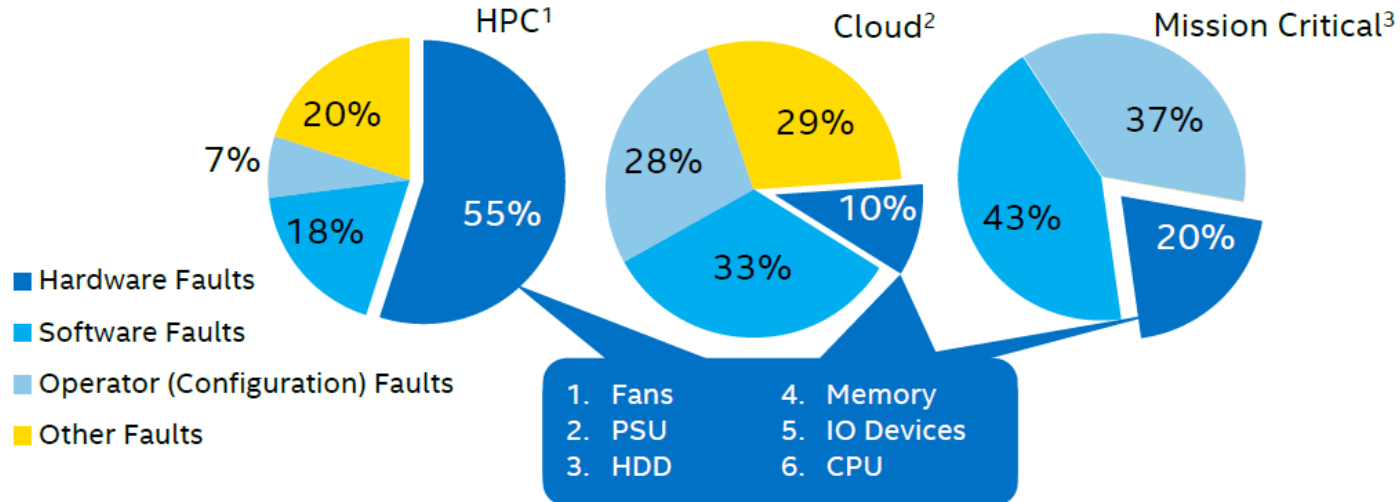
## How to prevent /minimize downtime



**Mission Critical Systems Require Advanced RAS features**

# RAS Needed in Mission Critical Systems

## Sources of service disruption

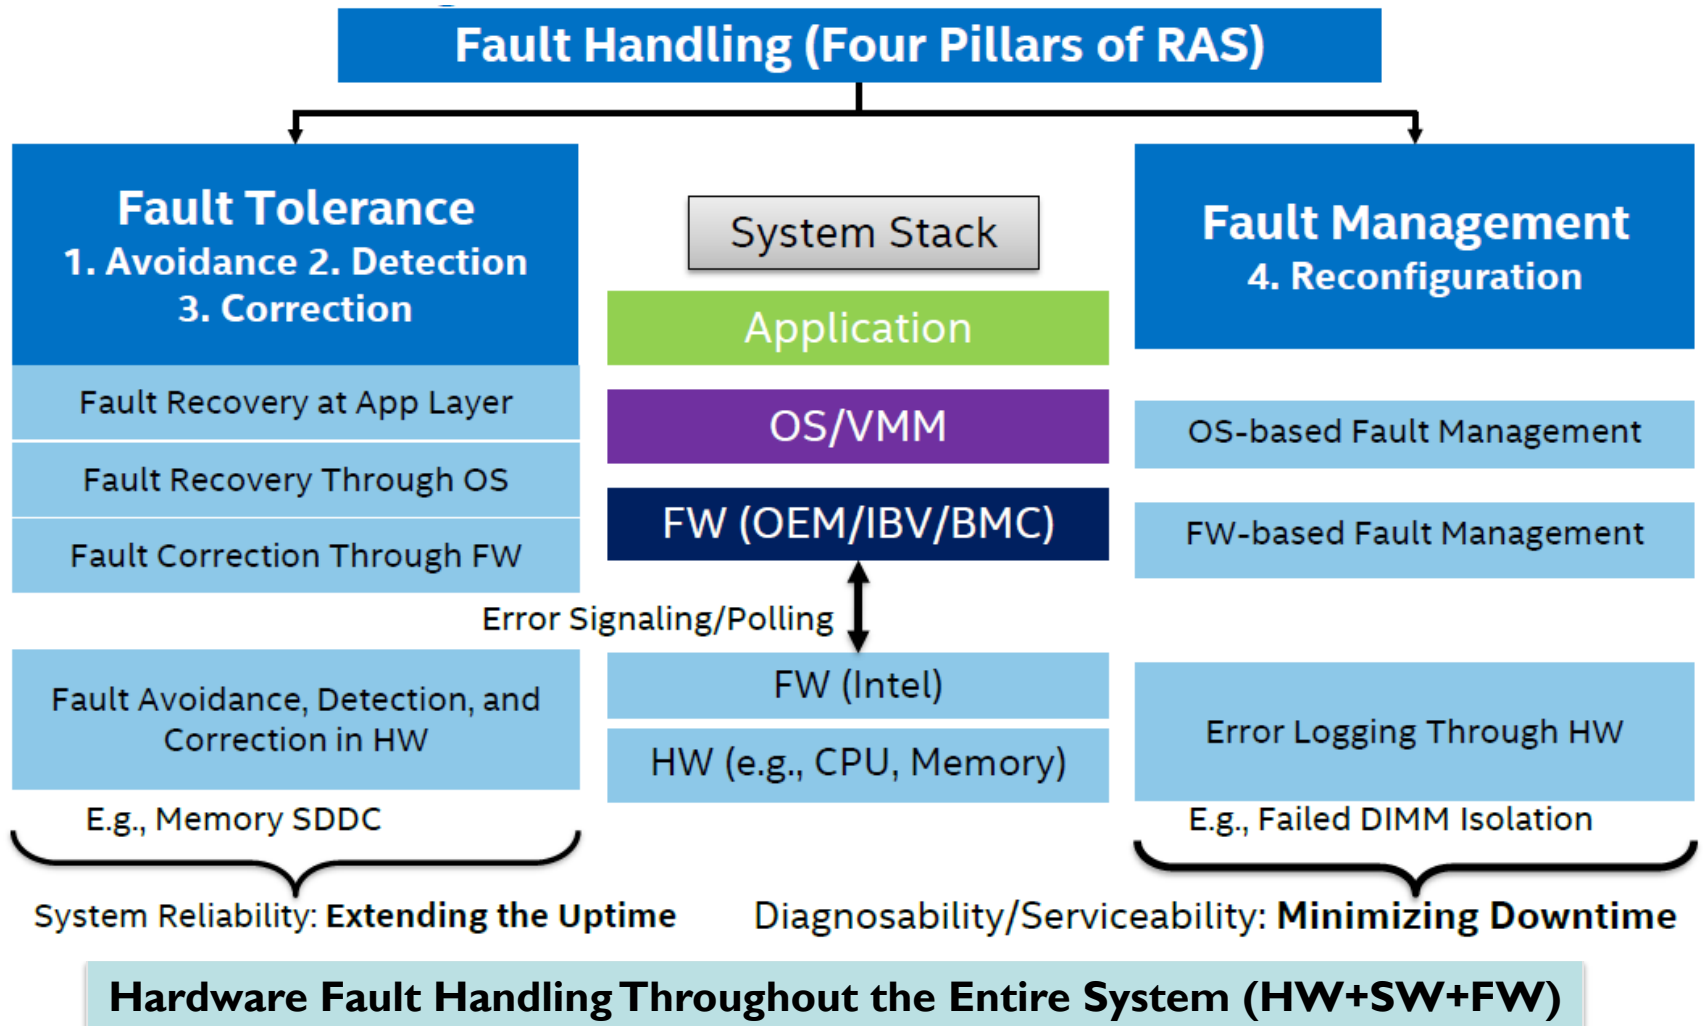


HW Fault Types	Definition
Soft Errors	High energy particle strike in caches and memory
Transient Errors	Electrical noise induced faults in links
Stuck-at Faults	Degradation of single memory cell, single lane on a link
Hard Failure	Failure of entire device (DRAM, memory Buffer, CPU)

### References:

1. Bianca Schroeder, Garth A. Gibson. (2007). Understanding failures in petascale computers. Journal of Physics: Conference Series, Vol.78, No. 1. (2007)
2. Urs Holzle and Luiz Andre Barroso. (2009). The Datacenter as a Computer, Google.
3. Best Practices for Continuous Application Availability, Gartner Data Center Conference 2008

# RAS Enabling Framework

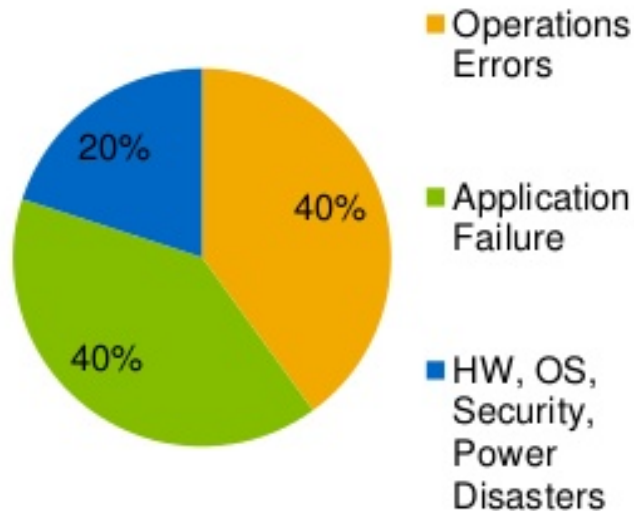




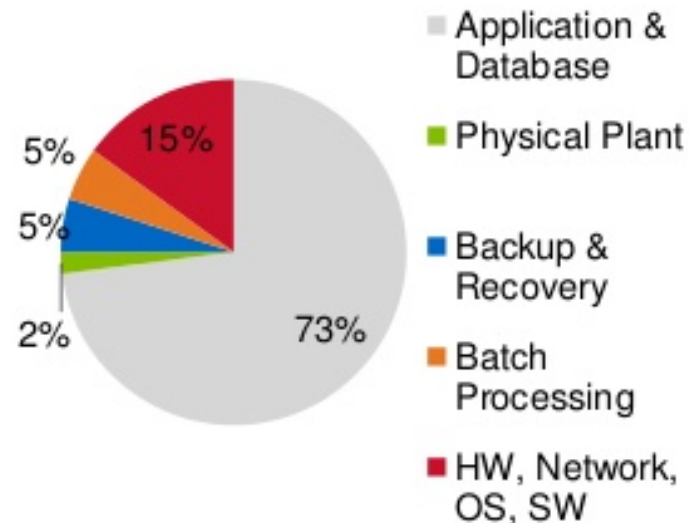
# RAS: Storage Systems

## Sources of Service Disruption

### Unplanned Outages



### Planned Outages



Source: Gartner 10/2011

**Application/Software responsible for 40% of Storage Server Failures**

# General RAS Features for Storage

- ❑ RAID – Storage Array Controllers (RAID – mirroring, striping, parity – Data Protection, Redundancy and Failover)
- ❑ ECC – Error Correcting Code implementation (Memory and In-flight data protection)
- ❑ Memory Scrubbing (Memory protection – prevents, corrects multi bit failures over time)
- ❑ DIMM sparing (Dynamic backup and Failover)
- ❑ Thermal Overload – prevention and mitigation
- ❑ Security Features

# RAS for Storage with NTB, ADR + Intel® QuickData Technology

- ❑ NTB = Non-Transparent Bridge
  - ❑ Dynamic memory data redundancy, failover, protection and security. Two or more nodes having mirror copies of the data yet blind to each other.
- ❑ ADR = Asynchronous DRAM Refresh
  - ❑ **PLATFORM STORAGE EXTENSIONS (PSE)** Data is protected in self-refreshed DRAM and/or copied to non-volatile storage (NVRAM, SSD or HDD) at the conclusion of the cycle.
- ❑ Intel® QuickData Technology
  - ❑ Allows rapid DMA of mirrored or lockstepped memory locations and NTB/PCIe endpoints across nodes without processor interrupts.

# Intel Server/Storage RAS Feature Summary

## System Reliability (Fault Tolerance)

ECC, Cache corruption containment

DRAM/Memory CMD/ADDR parity, Rank sparing, mirroring

Intel® UPI CRC, link data retry

PCI Express link retrain, CRC checking, ECRC

Virtual partitioning

## System Serviceability/Diagnosability (Fault Management)

Error Reporting, BIST, MCA Bank Error Reporting etc.

DDR4 CRC, Memory corrected error reporting

PCI Express “stop and scream”, card hotplug/hotswap

Failed DIMM isolation, enhance SMM, error injection capability

Color Key:

CPU
Memory
Intel® UPI
PCIe
System

# RAS Areas Covered

## NTB + Intel® QuickData and ADR Enhancements in RAS

Snapshots

Data Locality

Cloning/Mirroring

Data Tiering

Data Availability

Quality of Service (QoS)

Resilience

Compression

Erasure Coding

Deduplication

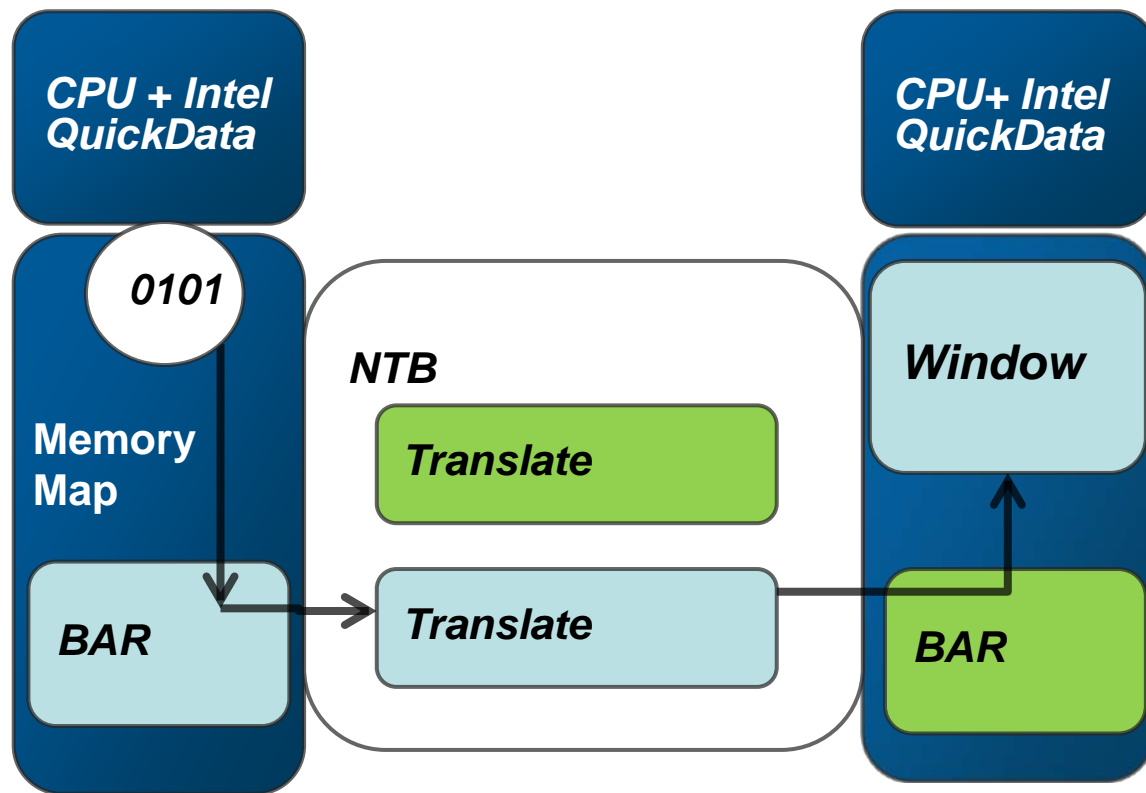
# RAS Enhancements for Storage with NTB

- ❑ **PCIe NTB – Non-Transparent Bridge Operation – Typical Storage Use Case:**
  - ❑ **Direct PCIe connection to another node over internal backplane for ideal redundancy and failover.**
  - ❑ **High bandwidth instant data interlink (via PCIe 3.0)**
  - ❑ **Maintain outstanding Quality of Service (QoS) for uninterrupted critical transaction with multi-node failover configuration.**
  - ❑ **Supports PCIe (3.0) Dual Cast\*\* - Maximizes memory channel bandwidth with the capability to perform a single write transaction to two targets with PCIe dual-cast**
  - ❑ **Allows host input/output controller (IOC) to write data directly to PCIe NTB and local memory**

14

# RAS Enhancements Storage with NTB

- NTB – Non-Transparent Bridge Operation opens high-bandwidth window into a remote duplicate storage node (redundancy and failover)



- “Window” is open to Remote Controller’s memory

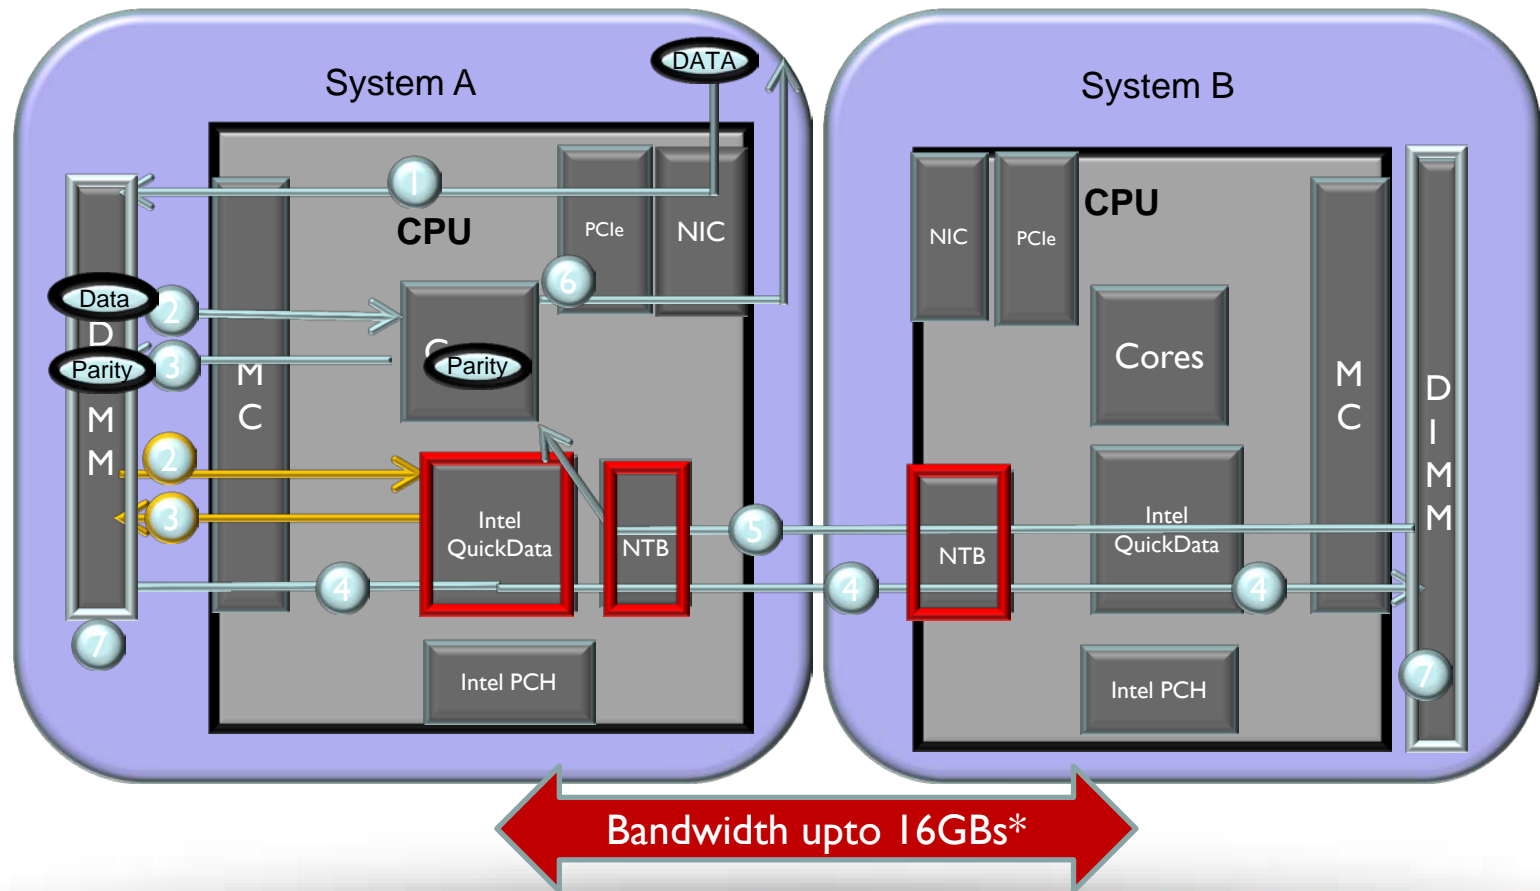
- BARs define Base, Size, and Limit of aperture
- BARs claim transactions and forward to peer
- Translate Registers define where transactions land in peer’s memory

**NTB Allows Fast Redundancy Over the PCIe NTB Port**

# RAS Enhancements Storage with NTB

- NTB – Non-Transparent Bridge Operation – Typical Storage Use Case
  - Direct memory access (DMA) to remote mirrored node (redundancy and failover)

## Typical Enterprise/Storage Use Case Data Flow



\*available on 5<sup>th</sup> Generation Intel® Xeon® Processor based platforms+



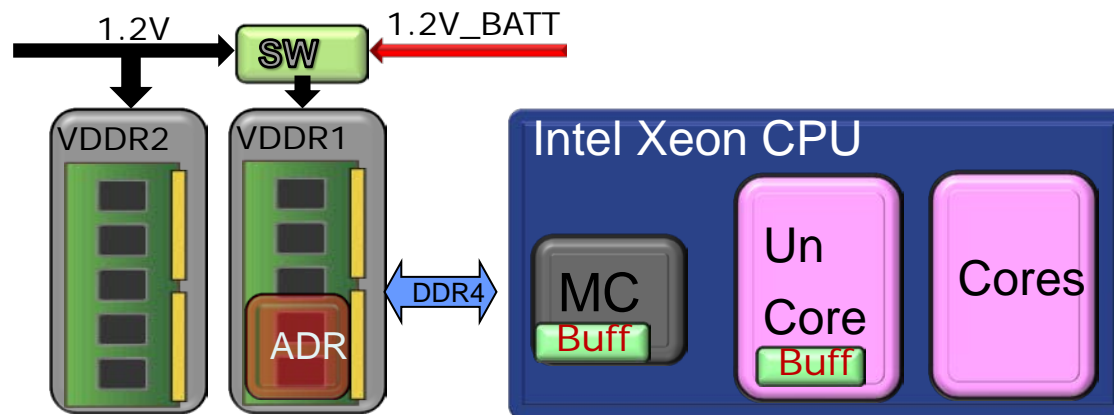
# Why Do I need NTB in a storage system?

**Problem: Need for fast backup node with real-time data mirroring.**

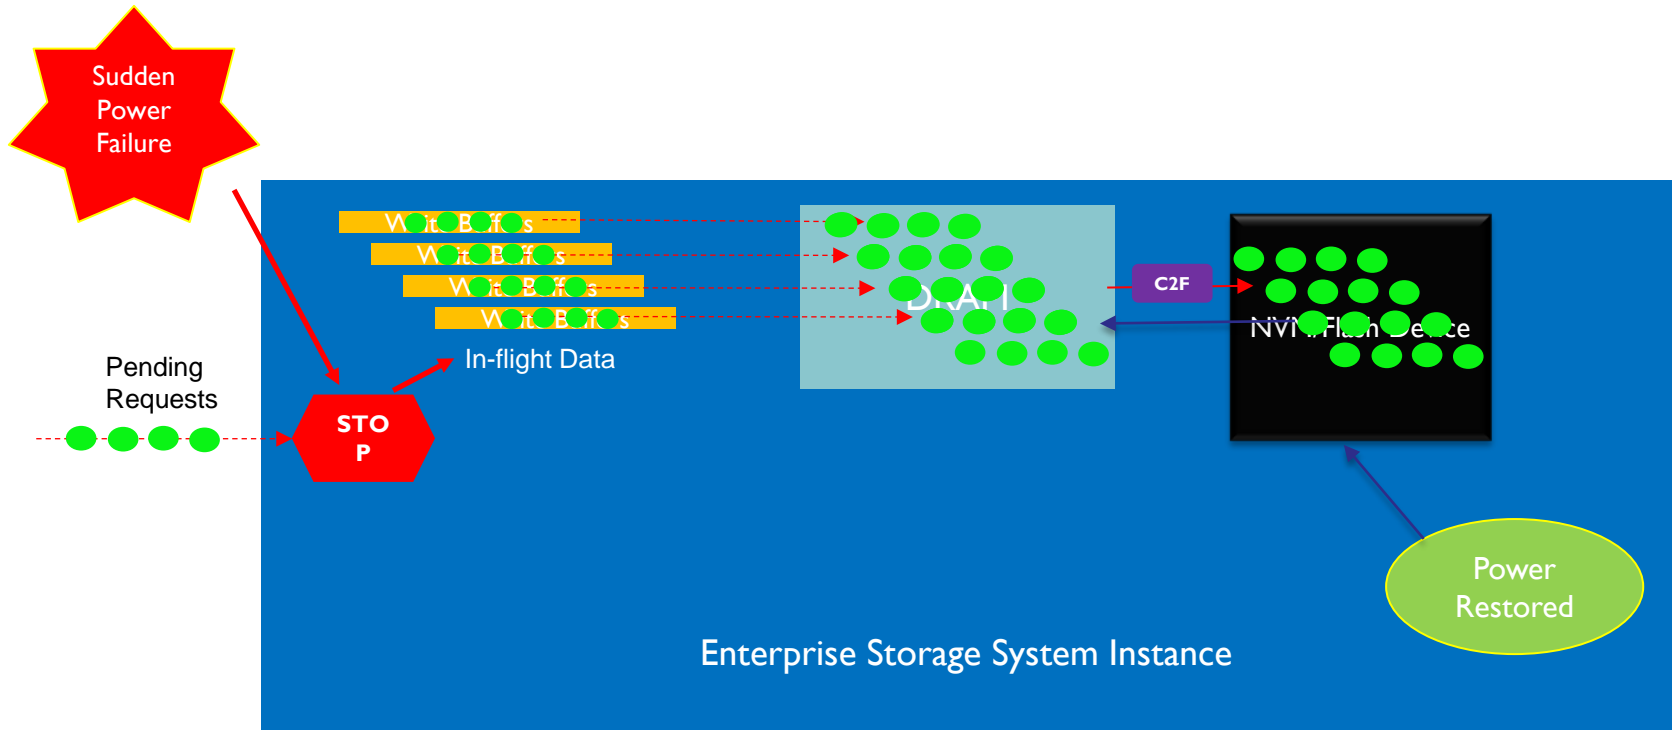
Solution	Implementation	Cost
Off-site/remote backup	Copy and duplicate all data to off-site or remote backup site.	<b>High</b> <b>(Capital Equipment)</b>  <b>High</b> <b>(Network Bandwidth)</b>
Non-Transparent Bridge (NTB)	Data is backed up seamlessly via fast PCIe interface without need to for any CPU computing resources over NTB link.	<b>Low</b> <b>(built-in)</b>  <b>No</b> <b>network bandwidth required.</b>

# Asynchronous DRAM Refresh (ADR)

- ❑ In the event of sudden power failure, ADR preserves key data in the battery backed DRAM or NVRAM.
- ❑ Example ADR sequence looks like the following:
  - ❑ Force flush transient data from protected 'write buffers' in the processor
  - ❑ Flush out to the battery DIMMs
  - ❑ Place DIMMs in battery backed self-refresh
  - ❑ DIMM content can be preserved by a backup battery or copying DIMM content to Flash or using a NVRAM.



# ADR in Action



# Why do I need ADR in a storage system?

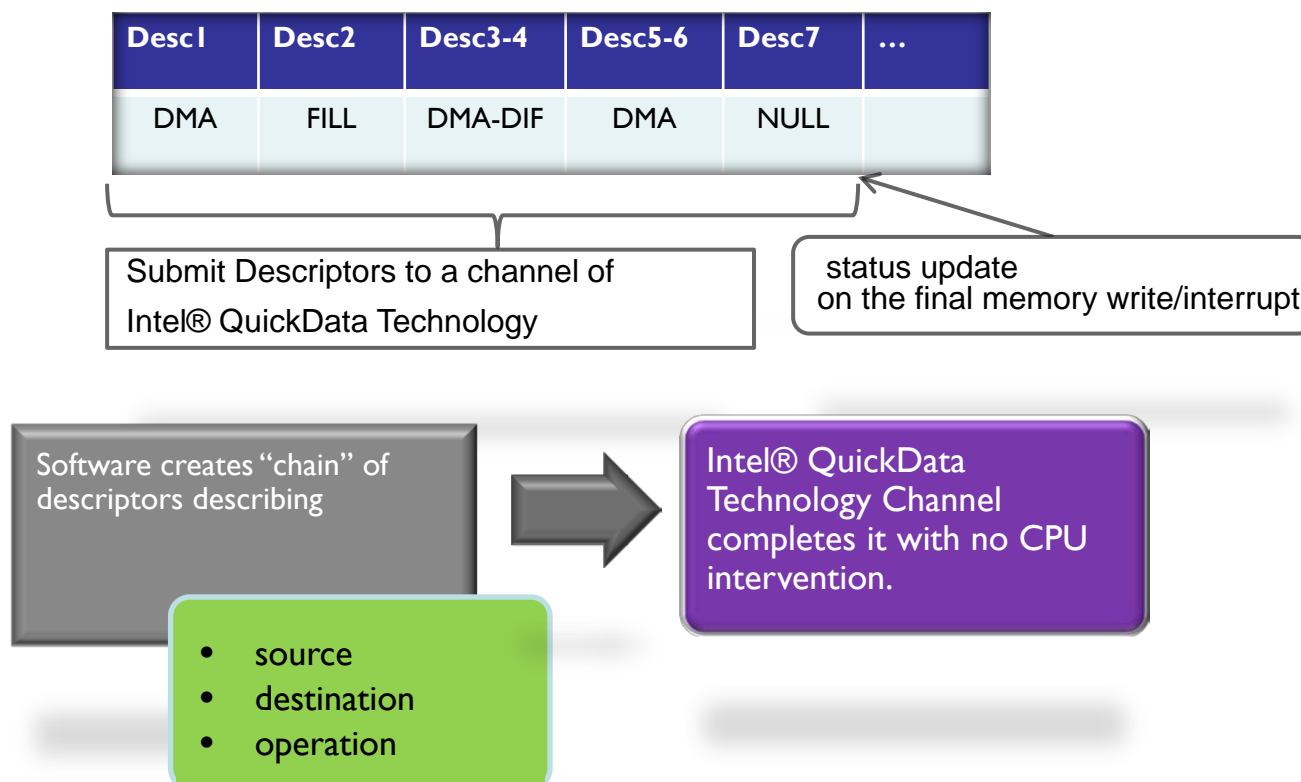
**Problem: Sudden Loss of AC Power, all transient data is lost.**

Solution	Implementation	Cost
Uninterrupted Power Supply (UPS)	In case of power failure – continuous, always on power for a limited period of time.	High
Remote Duplication/Off-site backup	Use real-time failover mechanisms that duplicate transient data real-time and transfer to a remote site for redundancy.	High
<b>Asynchronous DRAM Refresh (ADR)</b>	In the event of sudden power failure, all transient and in-flight data in local buffers is stored in battery backed memory or to persistent flash device.	Low (built-in)

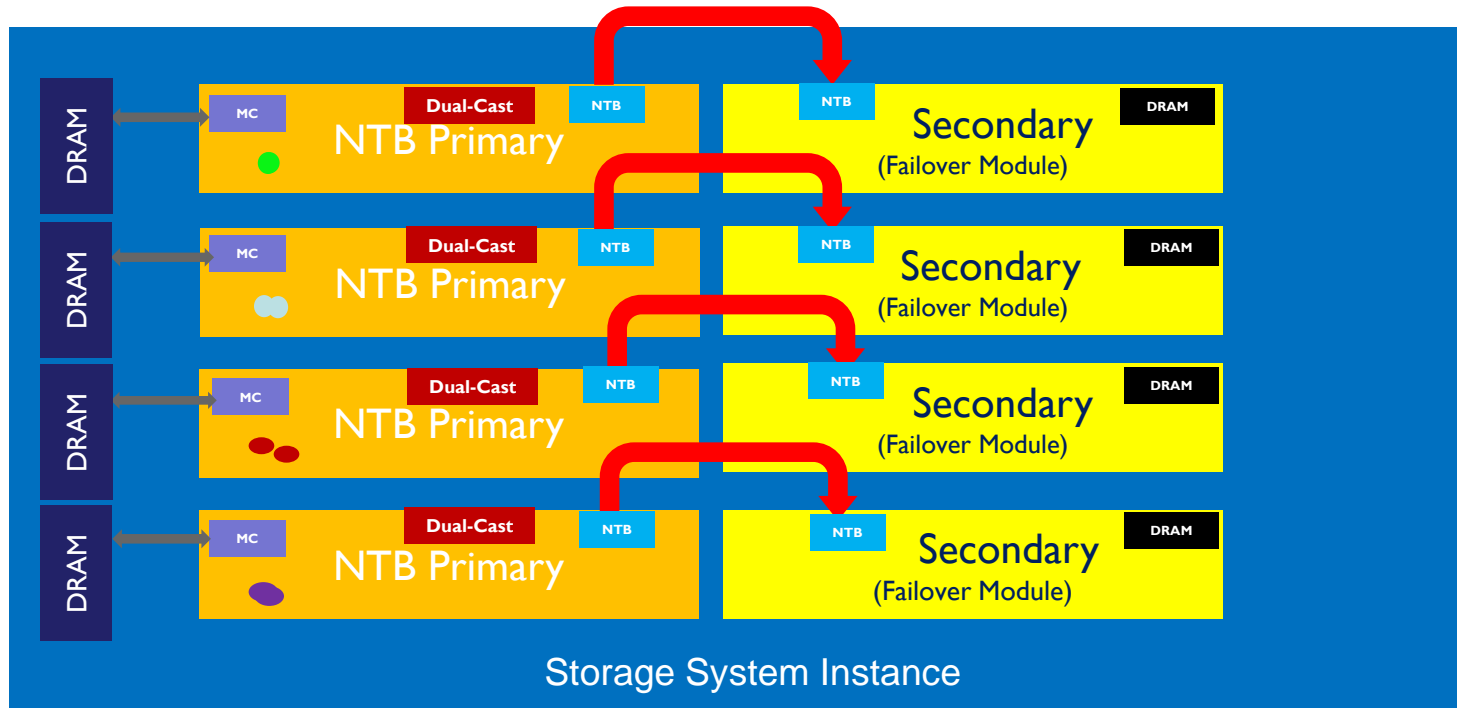
# High Speed Inter and Intra-node DMA Hardware Assist (Intel® QuickData Technology)

- ❑ Intel® QuickData Technology enables fast mirroring and lockstep of multi nodal systems and memory by concurrently duplicating data between primary and backup/failover nodes using a dedicated chipset hardware DMA engine while the processor continues to run other critical tasks.
- ❑ Intel® QuickData technology can work in conjunction with PCIe NTB (non-transparent bridge) to rapidly duplicate or mirror critical memory space regions.

# Intel® QuickData Technology Descriptors



# PCIe Dual-Cast + NTB + Intel® QuickData Implementation Example



**PCIe dual-casting allows simultaneous writes to DRAM and a secondary mirrored system over NTB**



# Why Do I Need Intel® QuickData Technology?

**Problem: Need to move large amounts of data concurrently to many nodes (and memory) at once without interrupting the processor cores.**

Solution	Implementation	Cost
Use CPU for data movement	Processor uses its resources to perform numerous calculations in order to move large data constantly between PCIe, I/O and Memory.	<b>High</b> <b>(Very CPU resource intensive)</b>
<b>Intel® QuickData Technology</b>	Intel QuickData technology moves data directly to/from memory and I/O devices offloading CPU resources completely. Also implements comprehensive RAS features like DIF and CRC error checking for data in-flight.	<b>Low</b> <b>(CPU not used)</b>



# RAS System Visibility

- ❑ **Processor**
  - ❑ **Machine Check Architecture (Error Logging Registers)**
    - ❑ IERR
    - ❑ MCERR
  
- ❑ **Platform Controller Hub (PCH)**
  - ❑ **Memory Error Corrections**
    - ❑ Parity
    - ❑ Scrubbing
    - ❑ SDDC
  - ❑ **Advanced Error Reporting (AER) – e.g. PCI Express**
    - ❑ ERR0
    - ❑ ERR1
    - ❑ ERR2
  - ❑ **End-to-End CRC (ECRC)**
  
- ❑ **Thermal Monitoring, Mitigation and Status**

# Final thoughts on Storage RAS

- ❑ **RAS features both at the platform as well as the system level are expanding rapidly with every generation. More granular visibility is being added into the systems in the next gen.**
- ❑ **Storage specific RAS features are also getting more expansive and granular over time – providing even more detailed look at the system level fault detection, recovery and prevention.**
- ❑ **With the growth of cloud storage and software defined storage – higher level RAS features for storage are becoming more predictive, preventative and self-healing in nature covering the entirety of the networked storage system.**
- ❑ **Intel's Platform Storage Extension (PSE) technologies have evolved significantly over time and remain the key drivers for providing comprehensive storage RAS features on Intel storage platforms.**