Zero-Maintenance Cache Protection Technical Brief

Reduced Data Center Operating Costs and Maximum Protection for Cached Data

July 2018

Microsemi
a Microchip company

Contents

Zero-Maintenance Cache Protection

Microsemi Adaptec 12 Gbps adapters with Zero-Maintenance Cache Protection (ZMCP) provide maximum protection for cached data and eliminate the substantial costs and environmental impact of complex, messy, and expensive lithium-ion batteries, while improving cache protection and performance. The following graphic illustrates how ZMCP works.

ZMCP

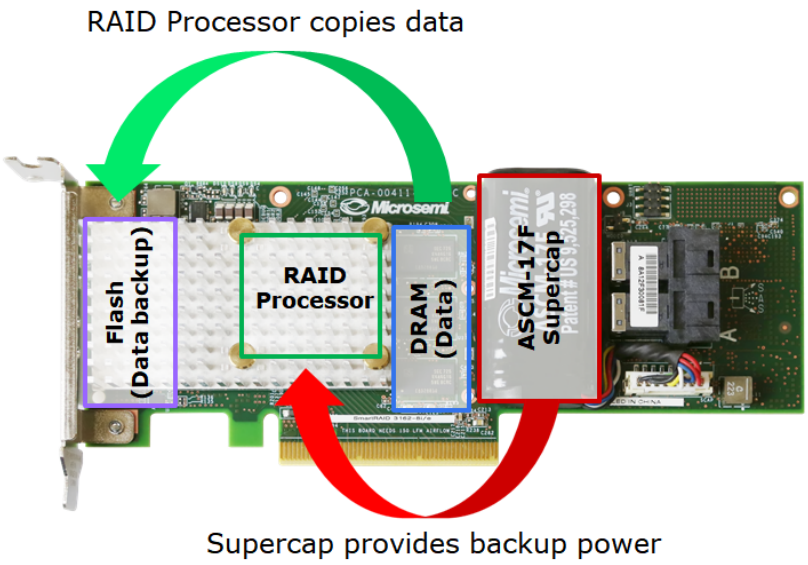

Features

- Low operating costs—No installation, monitoring, maintenance, disposal, or replacement costs due to batteries
- No data loss from power failures—Replaces lithium ion batteries
- Maintenance-free cached data protection—No need to monitor battery charge level, no shutdown required for battery replacement
- Protects data indefinitely—No rush to restart systems before the battery runs out, stores protected data for years
- Instant RAID cache protection—Charges in minutes instead of hours, RAID performance optimized immediately
- Environmentally conscious—No toxic battery disposal, simplified IATA compliance
- Flexible design—SmartRAID 315x includes integrated ZMCP. Third-generation ZMCP is available as an option for Series 8 adapters (AFM-700)
- SmartRAID 3162—First solution with fully integrated ZMCP, including onboard super capacitor to enable cache protection without requiring extra server space

The Need for Cache Protection

RAID 5 and 6 continue to be critical in enterprise storage systems as users continue to be concerned about data availability while optimizing capacity utilization for rapidly growing data sets. These RAID levels need some high-performance DRAM buffer to accelerate calculation of redundancy information (parity), that ideally is protected from data loss due to power failure. Additionally, in applications where spinning media (hard disk drives) are used, the use of DRAM caching can greatly enhance performance and reduce latency by up to 4x in real-world scenarios, even without the use of RAID. However, optimal performance can suffer unless the system operates with all available caches enabled, including write-back caching.

With adapter write cache enabled, data is stored in the adapter's memory and can be lost when a system power outage occurs. The traditional protection method in this scenario is a battery backup unit (BBU) installed directly on the adapter. This battery is used to maintain the data of the on-board memory cache until power can be resupplied to the unit.

Zero-Maintenance Cache Protection

Despite their obvious value, BBUs are not the optimal solution from a capital expense (CapEx) or an operating expense (OpEx) standpoint. BBUs require constant monitoring, maintenance, and replacement, and old batteries must be disposed of in an environmentally responsible manner.

SmartRAID 316x, SmartRAID 315x, and Series 8Q/8ZQ SAS/SATA RAID adapters offer a different approach, ZMCP.

The basic idea of ZMCP is to detect the loss of power to the adapter and then to copy the data in the onboard adapter cache to a non-volatile location—in this case, NAND flash memory of a type similar to that used in USB thumb drives and solid state disks. This process is supported by a super capacitor that keeps the necessary parts of the adapter active for the time required to perform the NAND flash copy. When power is finally restored to the adapter, the data in the flash memory is copied back to the onboard adapter cache and operation resumes as normal with all outstanding I/O requests preserved.

Once the data has been copied to the flash memory, the adapter no longer needs power in order to preserve the data.

Adaptec was first to deploy ZMCP with its 3 Gbps RAID adapters. The second generation of ZMCP was optimized for Adaptec 6 Gbps products. Each generation adds more advanced functionality and more integration of the capacitor module. The evolution of the advanced features is described in the next section. The sixth generation ZMCP, deployed on SmartRAID 316x, integrates the cache backup circuitry, NAND flash memory, and the field-replaceable super capacitor all onto the adapter itself, with no need for a tethered external super capacitor. With a tethered solution, the external super capacitor consumes space in the server, such as a neighboring PCIe slot. Saving space in the data center is a critical ingredient to deploying high-density data centers.

ZMCP Evolution

Microsemi most recently released its sixth generation of ZMCP-enabled products. This section describes the evolution of ZMCP in the Microsemi Adaptec 12 Gbps adapters.

Evolution of ZMCP

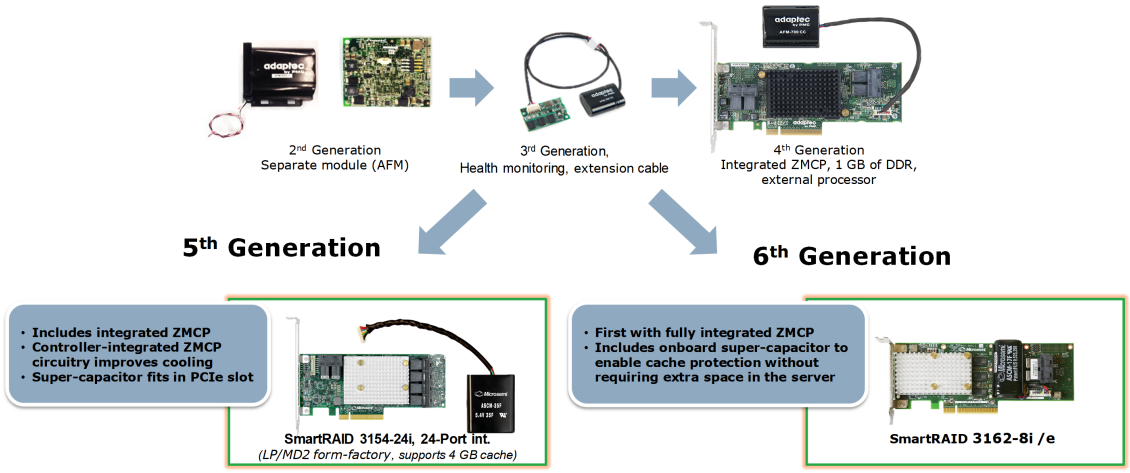

First and Second-Generation ZMCP

Recognizing the urgent need to protect cached data, Adaptec introduced the first-generation ZMCP solution on the SAS 3G SAS Adapter (Series 5Z) in 2009. It was the first PCIe SAS RAID adapter supporting ZMCP using a super capacitor. The second generation improved the first generation features and was deployed on Adaptec Series 6, the first SAS 6G RAID adapter series products.

Third-Generation ZMCP Advanced Features

The Series 8 and Series 8Q/8ZQ (12 Gbps) adapters used the AFM-700 flash-based cache protection module to enable instant battery-free cache protection. Unique to the Microsemi Adaptec 81605ZQ, the flash backup was embedded on the board, eliminating the need for a daughterboard. Third-generation ZMCP introduced real-time health monitoring, instant capacity-level monitoring, backup power to the adapter, and a new tightly integrated design.

Real-time Health Monitoring

Data center administrators could now instantly check the temperature, capacity, and remaining lifetime of the super capacitor through Microsemi Adaptec maxView, a web-based interface that made it simple to view, monitor, and configure all Microsemi Adaptec RAID adapters in a system. Real-time health monitoring was also available through the Microsemi Adaptec ARCCONF command line utility.

Instant Capacity Level Monitoring

ZMCP allowed capacity to be tested instantly without disrupting operations or impacting performance. In order to check the remaining capacity of a BBU, it must be discharged and then recharged, which can take up to 24 hours and leave cached data vulnerable in the event of a power failure.

Backup Power to the Adapter

The adapter logic provided the ability to monitor the loss of host power and switch to backup power from the AFM-700's super capacitors.

New Design

In this third generation, the ZMCP was tightly integrated into the adapter design with fewer parts that both reduced the footprint and achieved a higher mean time between failures (MTBF). Additionally, ZMCP in this generation used NAND flash in SLC mode (versus standard NAND flash used in previous versions), that allowed for better throughput, more reliable and faster transfers of backup cache data, and a longer product life cycle.

Fourth-Generation ZMCP Advanced Features

The Series 8Z adapters introduced the first integrated super capacitor solution for the Series 8 (offered on the 16 internal port adapter solutions only). Fourth-generation ZMCP lowered the cost for broad adoption of ZMCP with no separate module or PCB purchase required. The integrated ZMCP optimized RAID board cooling with no impact on airflow.

Fifth-Generation ZMCP Advanced Features for SmartRAID

The SmartRaid 315x includes integrated ZMCP on all boards with up to 24 ports and supports all Microsemi Adaptec internal and external adapter product SKUs. The SmartRAID 315x supports larger cache sizes (2 GB and 4 GB). Its super capacitor fits into a PCIe slot, and a mounting plate that ships with the RAID adapter can hold up to two adapters' super capacitor modules in one low-profile MD2 form factor PCIe slot. Also, new with this generation ZMCP is the capability to back up and protect up to 4 GB of DRAM cache on the adapter.

Sixth-Generation ZMCP Advanced Features for SmartRAID

The SmartRAID 3162-8i includes all the benefits of fifth-generation features plus onboard cache backup circuitry and flash memory. It integrates the ZMCP power source (capacitor module) in the adapter to enable instant cache protection without the need to find space to mount the super capacitor elsewhere in the system. The onboard capacitor module form factor fits into a PCIe slot and optimizes cooling/heat dissipation. It supports a five-year lifetime and is continuously monitored by the smart firmware to ensure the data can be safely backed up to the flash memory on the SmartRAID adapter.

SmartRAID 3162-8i

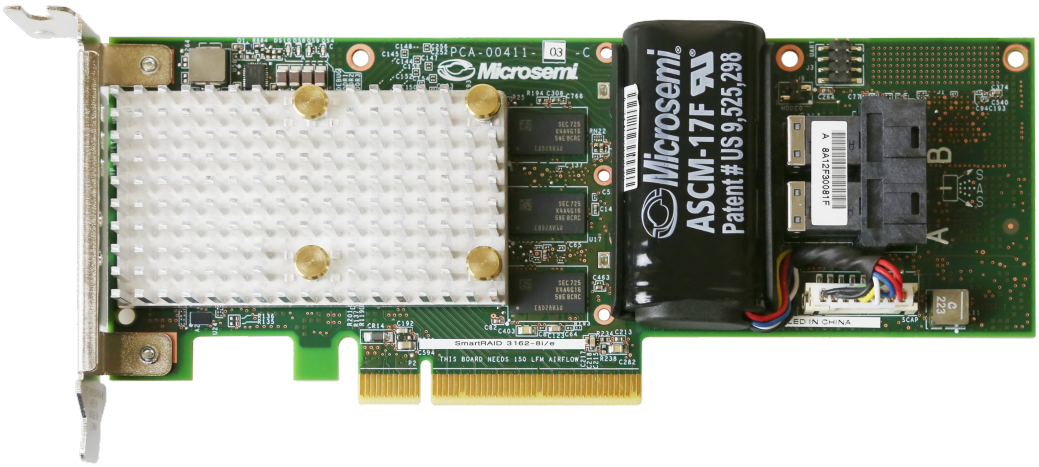

ZMCP Benefits Relative to BBUs

While BBUs have been an acceptable cached data protection solution for years, there are numerous hard costs, labor costs, and risk factors associated with managing and replacing BBUs after the initial purchase has been made.

Microsemi Adaptec RAID adapters with ZMCP eliminate all of these pitfalls. The following table summarizes the benefits of ZMCP.

ZMCP Benefits Relative to BBUs

Lithium-Ion Batteries	Consequences	Microsemi's ZMCP
Must be charged before they can be used	Cache is not enabled until the battery charge is complete	Charges instantly during system boot, enabling full protection
Must be "conditioned" during initial deployment	Add hours to the deployment process	No action required
Must be replaced on a regular maintenance cycle	Keep staff on hand to perform maintenance	No action required
Must be continually monitored so that failing batteries can be replaced	Add monitoring capability and corrective action process to operations process	No action required
Must be fully discharged then recharged in order to test capacity	Process can take 24 hours and leave cached data vulnerable to a power failure	Instant capacity monitoring with no disruptions
A failed battery must be replaced within 72 hours, perhaps less	Stock batteries in each location for urgent replacement	No action required
Must be correctly disposed of	Create, staff, and fund a disposal process	No action required

ZMCP Saves You Real Money

While the logic behind the ZMCP approach is relatively compelling, the actual financial impact is even more so. To compute these savings, we must look at the ways that people approach their existing BBU solutions.

Savings for the Meticulous BBU User

In this model, we assume the following about the way that the owner handles a card with a BBU.

- They purchase a new battery every year, as recommended, and keep a couple of replacements on hand all of the time to take care of unexpected events.
- They carefully schedule downtime for their users to change the battery. They attempt to deploy a replacement device to take up the load while the primary server is out of action. Several individuals in the IT department must usually coordinate their efforts to make this happen.
- Replacement batteries are allowed to fully charge before putting a system back into operation.
- Systems are continuously monitored to detect a failing battery.

In this case, we can assume that the chance of losing data due to a power failure happening while the battery is out of action is very small. The costs of this methodology include the following:

- Capital expense of purchasing batteries—one per year for four years, plus one extra to take care of unexpected issues.
- Operating IT costs to install the initial device, plan downtime, replace batteries, and recharge them.
- Potential "overtime" or "disruption" costs when a power failure occurs and systems have to be turned back on within the 72-hour battery charge window—potentially disrupting other activities, or occurring during the night, on a weekend, or over a holiday period. Even in the best case, the pressure to repower systems has an impact.
- Productivity impact on other users.

In an ideal world, the chance of losing data in the latter example would be zero due to all of the planning. However, in reality that is never the case. The following table illustrates the total cost of ownership (TCO) for a BBU solution.

Total TCO for BBU-Based Solution

Item per Server	Impact Over 4 Years	Cost per Item	4 Year Cost[1]
Batteries	4	$175 each	$700
IT manpower			
Initial installation	.5 hours	$20/hour	$10
Downtime preparation	1.5 staff-hours, 3 times	$30/hour	$135
Battery replacement	1 hours, 3 times	$20/hour	$60
Monitoring	30 seconds per day	$20/hour	$240
Productivity loss	5 people, 15% impact	$40/hour	$800
Activity cost of power loss	3 hours, 30% chance of occurring, 8 incidents	$50/hour	$360
Total			$2,305

1. Based on single server model.

Savings for the On-demand Replacement BBU User

The second model to consider is the BBU user who waits for the alert light to come on before doing anything.

To compute the impact of this scenario we assume the following.

- A replacement battery had been purchased at the time of the initial installation and is sitting on a shelf somewhere, ready to go. Another new battery is purchased to replace this one.
- The server is immediately taken out of service and, because this is an unplanned event, a relatively large number of people have their productivity affected.
- Because users are waiting to get back on the affected system, replacement batteries are NOT allowed to fully charge before putting the server back into operation. This minimizes the amount (and cost) of downtime for each replacement, but it exposes the system to potential data loss while the battery is charging. We assume the system takes two hours to fix, and users are allowed access after another three hours—a net total of five hours of lost work for each affected user.
- Systems are continuously monitored to detect a failing battery.

TCO for an On-demand BBU Solution

Item per Server	Impact Over 4 Years	Cost per Item	4-Year Cost
Batteries	4	$175 each	$700
IT manpower			
Initial installation	.5 hours	$20/hour	$10
Downtime preparation	0	$30/hour	$0
Battery replacement	2 hours, once	$20/hour	$40
Monitoring	30 seconds per day	$20/hour	$240
Productivity loss	40 people, 100% impact, 5 hours each incident	$40/hour	$8,000
Activity cost of power loss	3 hours, 30% chance of occurring, 8 incidents	$50/hour	$360
Lost business cost	2 hours of system downtime, $50M per year, 10% impact	$570/hour	$1,140
Total			$10,490

We have also included an impact on the business itself because this type of downtime is unscheduled. The cost of these instances is hard to compute, and estimates calculated by directly-impacted parties are sometimes as high as $500,000 per hour or more. We take a much more conservative approach and simply assume that the impacted device has a 10% impact on a business valued at $50 million per year.

The Bottom Line on Total Cost of Ownership (TCO)

Obviously, there are other scenarios that we could consider, but we have described two reasonable cases. The net result is:

Summary of TCO for a BBU Solution

Methodology	Cost Over 4-Year Lifetime
Taking proper care of your BBU	$2,305
Reacting to BBU emergencies	$10,490 (Plus data loss risk)
Zero Maintenance Cache Protection	Cost of the ZMCP module

Conclusion

By eliminating costly BBU technology and related expenses, ZMCP offers the industry's most complete and efficient data protection solution available today. In addition to immediate RAID performance optimization, Microsemi Adaptec RAID adapters with ZMCP are the environmentally conscious solution with no toxic battery disposal and simplified IATA compliance.

Zero-Maintenance Cache Protection is available for the products in the following table.

ZMCP Products

SmartRAID 316x	SmartRAID 315x	Series 8
SmartRAID-3162-8i with maxCache 4.0 with onboard cache backup and flash memory	SmartRAID-3154-24i with maxCache 4.0 and integrated cache backup	Series 8-81605ZQ with maxCache3.0— Flash backup circuitry built into the adapter (no need for daughter board)
SmartRAID-3162-8i /e with encryption and maxCache 4.0 with onboard cache backup and flash memory	SmartRAID-3154-16i with maxCache 4.0 and integrated cache backup	Series 8-81605Z— Flash backup circuitry built into the adapter (no need for daughter board)
	SmartRAID-3154-8i with maxCache 4.0 and integrated cache backup	Series 8-8885Q with maxCache3.0-AFM-700 included
	SmartRAID-3154-8e with maxCache 4.0 and integrated cache backup	Series 8-8885— AFM-700 sold separately
	SmartRAID-3152-8i with maxCache 4.0 and integrated cache backup	Series 8-8805— AFM-700 sold separately
	SmartRAID-3154-8i8e with maxCache 4.0 and integrated cache backup	Series 8-8405— AFM-700 sold separately
	SmartRAID-3154-8i16e with maxCache 4.0 and integrated cache backup	
	SmartRAID-3151-4i with maxCache 4.0 and integrated cache backup	

Microsemi Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

Microsemi, a wholly owned subsidiary of Microchip Technology Inc. (Nasdaq: MCHP), offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www microsemi.com.

ESC-2160437