

Series 7 in a Windows Environment

Microsemi Adaptec® RAID 72405 vs.

Competition: Illustrating the Performance

Benefits of High Native Port Count

White Paper



Introduction

Business and consumer demand for fast and reliable access to data and content continues to grow. Data centers are faced with the challenge of meeting that demand even as tight budgets force them to reduce the cost of service (COS) per user as well as hardware-related Capital Expenses (CapEx) and Operating Expenses (OpEx).

Compounding the challenge is the shrinking availability of physical space — many data centers are unwilling to take on the expense of more real estate and the associated costs that come from powering and cooling larger spaces.

Clearly, the days of simply adding more servers to accommodate more traffic are over. As a result, server vendors have introduced smaller, denser, server chassis sizes that help data centers add more storage I/O capability while maintaining the same or a smaller footprint.

Accordingly, storage vendors are focusing on small form-factor solutions that will fit into the smaller chassis while providing the high performance that data centers require. Storage components, such as hard disk drives, are getting physically smaller (though increasing in capacity). Low-profile storage adapters are becoming more common as well, but only a select few are able to deliver top I/O performance and low latency in such a compact form factor. This is an important consideration for data centers because finding the right combination of form, fit, and function allows them to deploy one SKU for almost any configuration and simplify everything from the purchase decision to installation to maintenance.

Series 7 SAS/SATA RAID adapters address the needs of space-challenged modern data centers by doubling storage performance compared to previous-generation RAID adapters while featuring high port counts and low-profile form factors that reduce latency, real estate requirements and power consumption.

The New Generation of PCIe

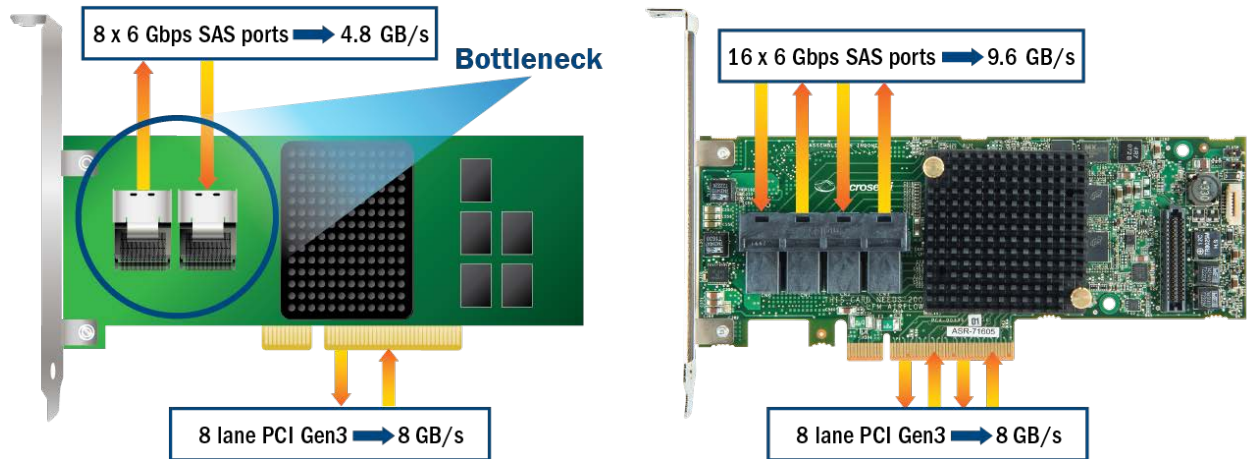
PCI Express (PCIe) is a motherboard-mounted expansion bus that, through a connected device such as a RAID adapter, connects the host system processor to add-on peripherals, such as storage systems. Introduced into servers and workstations in 2012, the third generation of PCIe (PCIe Gen3) doubles bandwidth to the host compared to its PCIe Gen2 predecessor, increasing per-lane throughput from 250 MB/s to 500 MB/s.

With PCIe Gen2, 8 6 Gbps SAS/SATA ports are sufficient to achieve maximum performance. However, PCIe Gen3 requires a minimum of 16 native 6 Gbps SAS/SATA ports to double the bandwidth through the storage connections.

A select group of storage adapters that claim to be designed for PCIe Gen3 has appeared the market, but most of them max out at 8 ports and cannot take full advantage of PCIe Gen3's superior performance.

Only Series 7 SAS/SATA RAID adapters, available with 16 or 24 native SAS/SATA ports, are designed to fully exploit the high-performance characteristics of PCIe Gen3, as shown in the following figure.

Figure 1 PCIe Gen3



Significance of High Native Port Count

In recent years, the storage industry has been transitioning from 3.5-inch storage drives to 2.5-inch Small Form Factor (SFF) drives as advancements in technology allow storage vendors to address the aforementioned physical space challenges faced by data centers. Not only do SFF drives offer the obvious advantage of allowing more drives to fit into the same server rack space, but 2.5-inch drives hold more capacity per space occupied than the 3.5-inch drives they are replacing. Indeed, SFF HDDs now boast storage capacities of 1 TB or more.

Additionally, the cost of 2.5-inch flash-based solid state drives (SSDs) is finally coming more in line with HDDs in terms of the traditional "cost per GB of capacity" metric. That, combined with a higher read bandwidth, higher input/output operations per second (IOPs), better mechanical reliability, and higher resistance to shock and vibrations compared to HDDs, is driving an industry-wide transition to SSDs.

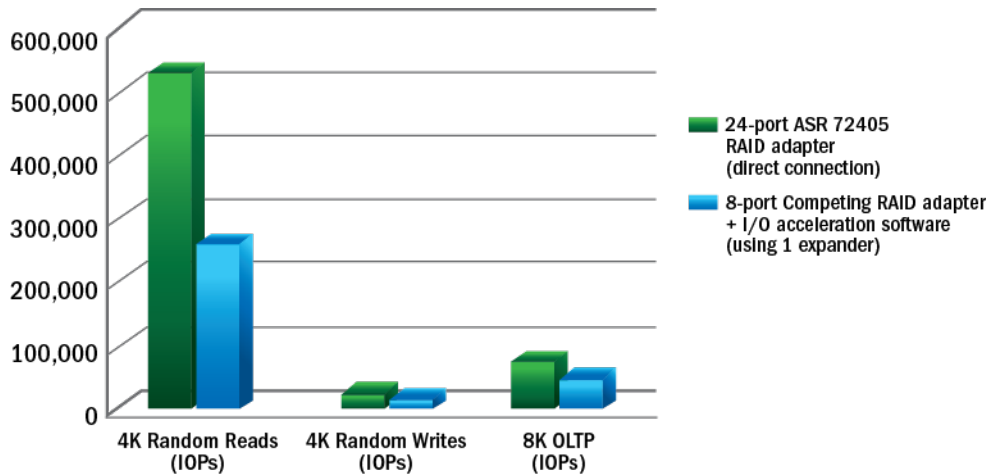
As the quantity of drives in a server chassis increases, the storage adapter card's port count requirements also increase.

The traditional method for increasing the storage adapter's port count has been through the use of an expander — a board that enables the connection of additional attached SAS or SATA devices when the adapter does not have enough ports to accommodate them. However, expanders have a number of limitations: not only do they add complexity, they also occasionally face compatibility issues with other components in the storage solution.

On top of that, expanders are notorious for causing latency and limiting data transfer bandwidth. Both of these issues have long been tolerated by data centers using HDDs, as they did not cause a huge impact on the already slow read and write speeds of HDDs. But as higher-performance SSDs gain traction in storage solutions, the latency and bandwidth issues of expanders have become more noticeable, and therefore less acceptable.

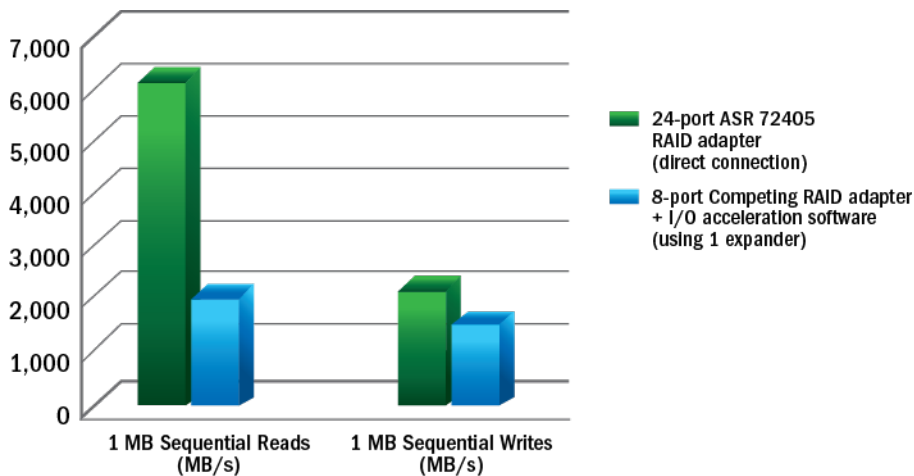
In a RAID 5 configuration using 24 SATA SSDs, the use of expanders causes a roughly 60% performance drop on random read IOPs, and a roughly 20% performance drop on OLTP read/write IOPs, compared to a direct connection through native ports. The following figure shows a RAID 5 configuration with 24 SATA SSDs.

Figure 2 RAID 5 Random Performance (24 SATA SSDs)



Similarly, in a RAID 5 configuration with SATA SSDs, as shown in the following figure, the use of expanders causes a roughly 70% performance drop on sequential read MB/s, and a roughly 40% performance drop on sequential write MB/s, compared to a direct connection through native ports.

Figure 3 RAID 5 Sequential Performance (24 SATA SSDs)



This problem can be partially overcome if SAS devices are used, since they are dual-ported and allow all 8 SAS port connections to be leveraged through the expander. However, as illustrated in the following figures, performance of the 8 6 Gbps SAS ports flattens out at the peak data rate and competing products cannot match Series 7 speeds.

Figure 4 RAID 5 Performance (24 SAS SSDs)

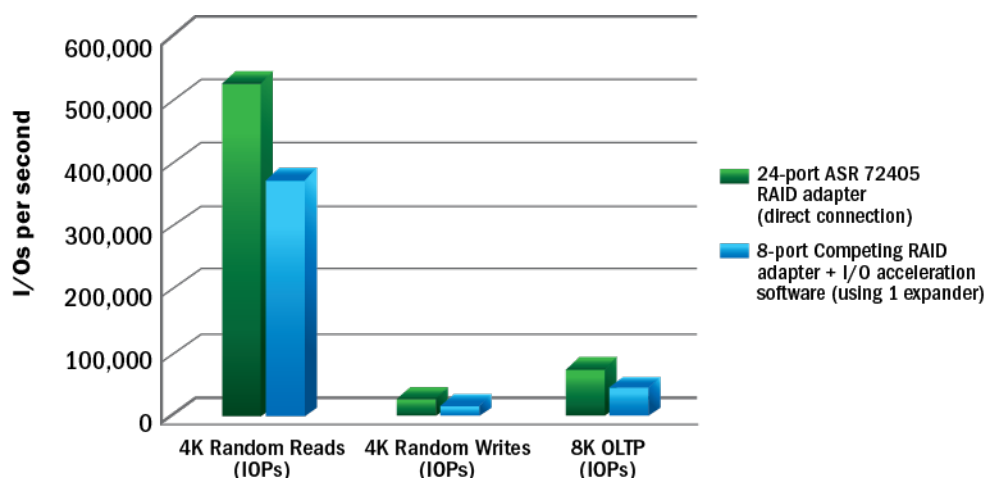
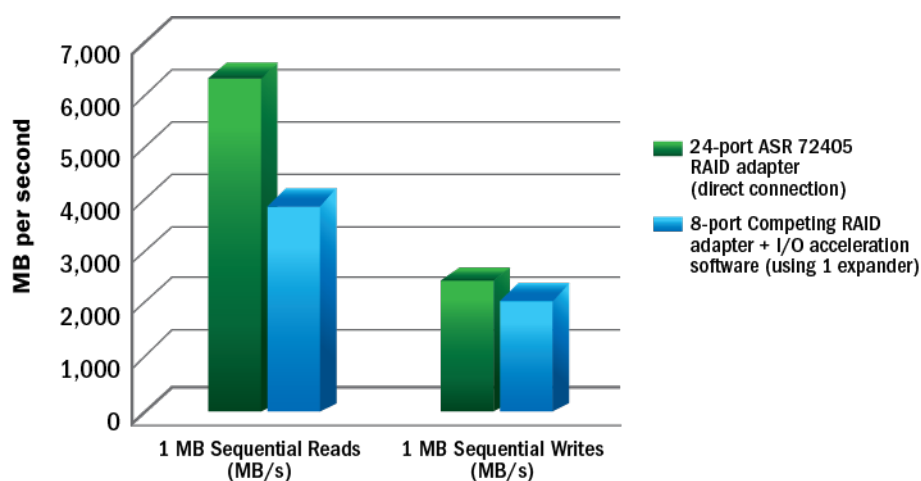


Figure 5 RAID 5 Sequential Performance (24 SAS SSDs)



Another drawback of expanders is the additional cost they add to a storage solution — about \$200 for the expander itself plus the cost of cables plus installation, increased power consumption, and maintenance costs.

An ideal solution for data centers would be a 6 Gbps storage controller with a high native port count that can take advantage of PCIe Gen3's performance.

However, as mentioned, most 6 Gbps storage adapters max out at only 8 ports.

Series 7

The Series 7 SAS/SATA RAID adapter family features Microsemi's 24 port PM8015 RAID-on-Chip (ROC), which combines an x8 PCIe Gen3 interface with 6 Gbps SAS ports to enable a new generation

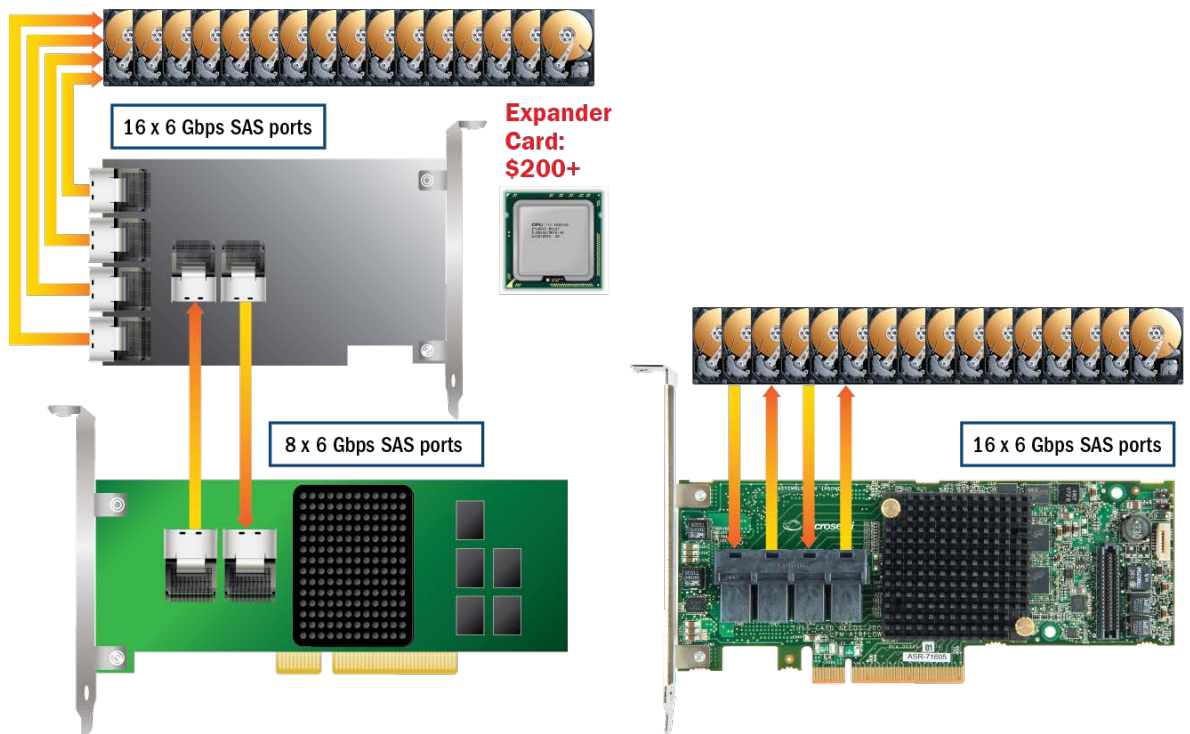
of high performance, high native port count RAID adapters that are unmatched by any other ROC in the industry.

Traditionally, RAID adapter performance has centered around read and write throughput, measured in megabytes per second (MB/s). Using this metric, Series 7 adapters perform up to 83% better than competing RAID adapters — 6.6 GB/s on sequential reads and up to 2.6 GB/s on sequential writes on parity RAID 5.

Moreover, with the popularity and growth of SSDs, input/output operations per second (IOPs) is emerging as the new “lead horse” in performance metrics, with the most common configuration being the 4K random-read number. Using 4K I/O size in random scenarios is driven by the fact that most operating systems use 4K cache sizes in the server DRAM and, with that, 4K is typically the smallest I/O size for random workloads. In a RAID 5 configuration with 16 direct-connected SSDs, Series 7 adapters again lead the field with 450K IOPs — nearly 10x the performance of previous-generation RAID adapters, and more than double that of the competition.

As noted earlier, RAID adapters with only 8 native ports cannot pass PCIe Gen3’s performance gains through from the bus to the storage connections. Series 7 RAID adapters are the first on the market to take full advantage of PCIe Gen3 performance gains by using HD mini-SAS connectors to offer options with 16 or 24 native SAS/SATA ports, as shown in the following figure.

Figure 6 Configuration Complexities and Costs Expanders vs. Direct Connect



Performance Comparison

The following data compares the 24-port PCIe Gen3 ASR 72405 SAS/SATA RAID adapter to a competing 8-port PCIe Gen3 RAID adapter.

72405 Performance Summary:

- Max Read IOPs observed: 534K IOPs
(RAID-0: 24x SSD: Random Read 4 kB: Queue Depth = 64)
- Max Write IOPs recorded: 374K IOPs
(RAID-0: 24x SSD: Random Write 4 kB: Queue Depth = 64)
- Max Read throughput observed: 6400 Mbps
(RAID-0: 24x SSD: Sequential Read 1 MB: Queue Depth = 64)
- Max Write throughput observed: 5229 Mbps
(RAID-0: 24x SSD: Sequential Write 1 MB: Queue Depth = 256)
- Coalesce Read IOPs performance is observed at: 1.02M IOPs
(RAID-0: 24x SATA Sequential Read 512 B: QD=256)
- Coalesce Write IOPs performance is observed at: 1.07M IOPs
(RAID-0: 24x SATA Sequential Write 512 B: QD=256)

Table 1 Specifications

Platform:	
Intel CanoePass Qty. 1x Intel Xeon 2.9 Ghz (8 Cores) (LGA2011) 2x 4 GB (or 8 GB total) DDR3-1333 Mhz ECC	
HBA:	
Cache mode	SSD: Logical Read=OFF / Write=OFF / Physical Write=ON SAS/SATA: Logical Read=ON / Write=ON / Physical Write=ON
ASR 72405	Fw B20104/Drv B20100
Competing adapter ¹	Fw v23.7.0-0035 / Drv v5.2.127.64 / MSM 12.05.03.00
Enclosure	
Supermicro SAS216A, 24-bay SAS 6.0 Gbps Passive Backplane Supermicro SAS2-216EL2, two independent 24-bay Dual SAS 6.0 Gbps Expander Backplanes (Dual SMC expander backplanes needed to avoid SATA performance bottleneck)	
Device: Qty 24x	
SSD	OCZ Deneva "R" Series SLC: D2RSTK251S14-005 6 Gb 50 GB Fw3.00E
SAS	Seagate Savvio 15K.3: ST9146853SS 6 Gb 146 GB Fw0002
SATA	Seagate Constellation.2: ST9500620NS 6 Gb 500 GB FwSN02
Software	
OS	Windows 2008 R2 SP1 64 bit
IOMETER	Version: 2006.07.27 Workers: 4 Workload: SeqRead 512 B/1 MB, SeqWrite 512 B/1 MB Workload: RndRead 512 B/4 kB, RndWrite 512 B/4 kB Workload: OLTP 4 kB/8 kB Ramp/Run: 10 sec. / 30 sec.

1. I/O acceleration software, costing additional ~ \$150, enabled through using software key module

The following data compares the 24-port PCIe Gen3 ASR 72405 SAS/SATA RAID adapter to a competing 8-port PCIe Gen3 RAID adapter.

Table 2 RAID 5 Performance Comparison

RAID – 5 SSD	Based on Queue Depth = 256				
	24-port ASR 72405 (direct connection)	Competing 8-port adapter (No I/O acceleration software)	%	Competing 8-port adapter (with I/O acceleration software)	%
512 B Random Reads (IOPs)	531,580	163,283	225.6%	373,054	42.5%
512 B Random Writes (IOPs)	28,311	16,771	68.8%	16,759	68.9%
4K Random Reads (IOPs)	529,731	163,283	224.4%	377,448	40.3%
4K Random Writes (IOPs)	28,488	16,655	71.1%	16,660	71.0%
1 MB Sequential Reads (MB/s)	6,192	3,928	57.6%	3,904	58.6%
1 MB Sequential Writes (MB/s)	2,237	2,185	2.4%	2,190	2.2%

Note: “%” references the comparison of ASR 72405 over competing product.

Table 3 RAID 10 Performance Comparison

RAID – 10 SSD	Based on Queue Depth = 256				
	24-port ASR 72405 (Direct connection)	Competing 8-port adapter (No I/O acceleration software)	%	Competing 8-port adapter (with I/O acceleration software)	%
512 B Random Reads (IOPs)	534,824	170,338	214.0%	373,396	43.2%
512 B Random Writes (IOPs)	130,851	43,901	198.1%	43,752	199.1%
4K Random Reads (IOPs)	535,546	170,092	214.9%	364,165	47.1%
4K Random Writes (IOPs)	192,051	43,808	338.4%	43,508	341.4%
1 MB Sequential Reads (MB/s)	5,433	3,911	38.9%	3,888	39.7%
1 MB Sequential Writes (MB/s)	2,892	1,418	104.0%	1,404	106.0%

Note: “%” references the comparison of ASR 72405 over competing product.

Table 4 RAID 0 Performance Comparison

RAID – 0 SSD	Based on Queue Depth = 256				
	24-port ASR 72405 (Direct connection)	Competing 8-port adapter (No I/O acceleration software)	%	Competing 8-port adapter (with I/O acceleration software)	%
512 B Random Reads (IOPs)	531,852	169,822	213.2%	375,676	41.6%
512 B Random Writes (IOPs)	238,681	166,875	43.0%	352,385	–32.3%
4K Random Reads (IOPs)	527,425	170,024	210.2%	377,964	39.5%
4K Random Writes (IOPs)	245,327	166,803	47.1%	295,252	–16.9%
1 MB Sequential Reads (MB/s)	6,182	3,467	78.3%	3,435	80.0%
1 MB Sequential Writes (MB/s)	5,229	2,790	87.5%	2,751	90.1%

Note: “%” references the comparison of ASR 72405 over competing product.

Conclusion

In order to continue meeting customer demand for fast and reliable access to data and content, data centers must employ efficient storage solutions that maximize I/O capability while fitting within budgetary and physical space requirements.

A new generation of PCIe Gen3 storage adapters seek to enhance storage I/O performance, but only Series 7 SAS/SATA RAID adapters fully address the needs of space-challenged modern data centers by offering 16 ports required to maximize PCIe Gen3 performance.

Series 7 adapters perform up to 83% better than competing RAID adapters in read and write throughput — 6.6 GB/s on sequential reads and up to 2.6 GB/s on sequential writes on parity RAID 5 — and lead the field with 450K IOPs — nearly 10x the performance of previous-generation RAID adapters, and more than double that of the competition.

**Microsemi Corporate Headquarters**

One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113

Outside the USA: +1 (949) 380-6100

Fax: +1 (949) 215-4996

Email: sales.support@microsemi.com

www.microsemi.com

©2016 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are registered trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

About Microsemi

Microsemi Corporation (Nasdaq: MSCC) 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-2161355