

SCHOONER TECHNOLOGY WHITE PAPER

Deploying Higher Level Building Blocks for Web 2.0 and Cloud Computing Datacenters

How Tightly Coupled Data Access Appliances Simplify Scaling,
Decrease Business Complexity, and Cut TCO

Table of Contents

Introduction	3
Computing Trends and New Technologies.....	4
Multi-Core Processors	4
Flash Memory.....	5
The Intel® X25-E Extreme SSD	7
Incorporating Flash Memory into Overall System Architecture	7
High-Performance Interconnects	7
Loosely Coupled Scale-Out Architectures.....	8
Challenges of Utilizing New Technologies	9
The Schooner Solution: Tightly Coupled Building Blocks.....	10
Scalable Hardware Platform	10
Schooner Operating Environment.....	11
MySQL and Memcached Appliances	12
Schooner Administrative and Analysis Tools	16
Benefits of Schooner Appliances	17
About Schooner Information Technology	17

Introduction

Today's Web and cloud computing datacenters have reached a critical juncture as exploding demand for their services has collided with existing architectures and technologies. These datacenters leverage standard low-cost X86 servers, Gigabit Ethernet interconnect, and open-source software to build scale-out applications with tiering, data, and application partitioning, dynamic random access memory (DRAM)-based content caching servers, and application layer node failure tolerance.

These loosely coupled architectures have enabled service scaling—but at a very high cost. Today's datacenters are reeling from the high costs of power, capital equipment, network connectivity, and space, and are hindered by serious performance, scalability, and application complexity issues.

Tremendous technology advances have appeared in recent years. Multi-core processors, flash memory, and low-latency interconnects offer tremendous potential improvements in performance and power at the component level. But when Web 2.0 and cloud computing datacenters integrate these new component technologies, they are severely underutilized. Serious scaling, performance, power, network connectivity, space, and complexity issues still remain.

Adapting an application and putting in place an optimized operating environment and the necessary technology configurations to effectively take advantage of multi-core processing, flash memory, and high performance interconnects requires major engineering and research efforts. The functionality required to achieve the benefits of these new technologies includes:

- High thread-level parallelism at all levels
- Granular concurrency control at each level
- Data/thread affinity management down to the core level
- Optimized path lengths for data access and thread context switching
- Multiple caching levels to exploit component technology access time and bandwidth variations
- Optimization for the many idiosyncrasies of simultaneous multi-threading and flash memory
- Selecting, tuning, and balancing component technologies and configurations to match workloads
- Measurement and analysis tools for implementation and deployment optimization
- Application and storage distribution
- Failure tolerant replication and recovery mechanisms

Web 2.0 and cloud computing enterprises must focus all resources on their core business of providing leading-edge application services. Higher level building blocks are needed that can effectively exploit these advanced technologies, to fundamentally solve the problems of power, performance, network connectivity, space, and complexity.

Scientists and engineers at Schooner Information Technology have succeeded in addressing all of these challenges. The results are tightly coupled, compatible, scalable data access appliances—yielding order-of-magnitude improvements in performance, power, network connectivity, and space at the datacenter level for cost-effective service deployment.

This white paper begins with a discussion of recent computing trends, taking a detailed look at new processor, flash memory, and interconnect technologies. Next, the issues regarding loosely coupled applications as they relate to exploiting these technologies will be examined. The benefits and requirements of tight coupling will be presented, followed by technical details of the Schooner appliance architecture and design. Finally, the paper will introduce the Schooner data access appliances and the transformational results they deliver for Web 2.0 and cloud computing datacenters.

Computing Trends and New Technologies

Processor, memory, and interconnect technologies have advanced dramatically over the last decade. Multi-core processors, flash memory, and high speed interconnects offering order-of-magnitude improvements in performance, power, network connectivity, and space are now standard and commoditized. But when Web 2.0 and cloud computing datacenter solutions try to integrate these technologies, they are severely underutilized—and performance, power, network connectivity, and space are still major issues.

To address these challenges, holistic system architectures that match workload, operating environments, and hardware organization are required. The result is deeply integrated, scalable building blocks that effectively harness the capabilities of these powerful technology advances to meet the key needs of Web 2.0 and cloud computing datacenters.

Multi-Core Processors

Multi-core processors place many processors and shared caches on a single chip, providing very high potential performance throughput for workloads with thread level parallelism. The industry-leading example of multi-core processor technology is the next generation [Intel® Core™ i7 processor](#), code-named Nehalem.




Figure 1. The Intel Core i7 processor, code-named Nehalem

Intel Nehalem processors provide 16 simultaneous hardware threads with shared level 3 caches on a single chip, directly connected to high speed memory. Multiple Intel Nehalem chips can be connected on a server board with all of their caches being coherent.

To fully realize the benefits of advanced multi-core processors, applications and operating environments need to have many parallel threads with very fast switching between them. They also need to support memory affinity and have granular concurrency control to prevent serialization effects.

Flash Memory

Flash memory is a non-volatile computer memory that can be electrically erased and reprogrammed. Flash memory has many promising characteristics—but also many idiosyncrasies.

Flash memory offers access times that are 100x faster than those of hard disk drives (HDDs), and requires much less space and power than HDDs. It consumes only 1/100th the power of DRAM, and can be packed much more densely—providing much higher capacities than DRAM. Flash memory is also cheaper than DRAM and is persistent when written, whereas DRAM loses its content when the power is turned off. Flash memory can be organized into modules of different capacities, form factors, and physical and programmatic interfaces.

However, flash memory access times are much slower than DRAM, and flash memory chips have write access behavior that is very different than their read access behavior. Flash memory writes can only be done in large blocks (~128 kB), and before writing, the region needs to be erased. Also, flash memory has limits on how many times it can be erased (~100k). As a result, small writes need to be buffered and combined into large blocks before writing (write coalescing), and block writes need to be spread uniformly across the total flash memory subsystem to maximize the effective lifetime (wear leveling).

The latency, bandwidth, capacity, and persistence benefits of flash memory are compelling. However, effectively incorporating flash memory into system architectures requires specific design and optimization—starting at the application layer, throughout the operating environment, and down to the physical machine organization.

All Flash Memory Is Not the Same

Flash memory chips are constructed from different types of cells (NOR and NAND), and with different numbers of cells per memory location (single-level cell or SLC; and multi-level cell or MLC). These variations result in very different performance, cost, and reliability characteristics. Table 1 summarizes the characteristics of the variations of flash memory as compared to HDDs and DRAM:

	Read B/W	Write B/W	Erase Lat.	Read Lat.	Cost per GB
HDD	100 mb/s	150.00 mb/s		5,000.00 us	\$0.10
NAND MLC	250 mb/s	70.00 mb/s	3.5 ms	85.00 us	\$3.50
NAND SLC	250 mb/s	170.00 mb/s	1.5 ms	75.00 us	\$11.00
NOR SLC	58 mb/s	0.13 mb/s	5,000.00 ms	0.27 us	\$70.00
DRAM	2,000 mb/s	2,000.00 mb/s		0.08 us	\$75.00

Table 1. Characteristics of the variations of flash memory as compared to disk and DRAM.

NOR vs. NAND

NOR flash memory chips have much lower density, much lower bandwidth, much longer write and erase latencies, and much higher cost than NAND flash memory chips. For these reasons, NOR

flash has minimal penetration in enterprise deployments; rather it is primarily used in consumer devices. Leading solid state drives (SSDs) are all designed with NAND flash.

NOR flash is much less dense than NAND flash; NAND flash memory chips have 2-8x the capacity of an equivalent size NOR flash memory chip. NAND flash has 5x the bandwidth of NOR flash for large reads, and 100x the write bandwidth of NOR flash. NOR flash also has poor erase bandwidth. When erasing any data on a NOR flash memory chip, the entire chip is tied up for half a second, blocking access to any other parts of the chip.

NOR flash's major benefit is that small reads can be performed more quickly than NAND flash, which makes it useful for certain types of read-mostly applications, such as cell phones. NOR flash has minimal investment by tier-1 memory suppliers, and very few manufacturing sources.

Flash: SLC vs. MLC

Another distinction with flash memory is SLC and MLC. MLC increases density by storing more than a single bit per memory cell. With increased density, MLC cost is lower than SLC, however the MLC write bandwidth and erase latency is about 2.5 times slower than SLC. In addition, MLC has a significantly lower lifetime than SLC.

Flash Interface: SSDs vs. PCIe

The most effective way to install flash memory into a server is using a SSD. SSDs interface through the SATA or SAS interface through which one would normally attach a HDD. Flash memory can also be installed into a server through interface cards that are plugged directly into the PCIe slots on the server motherboard (PCIe flash).

The flash memory management functions of write coalescing, space management, mapping, and wear leveling require significant on-going computation. PCIe-based flash memory subsystems perform these functions using the server's processor cores, consuming a large amount of CPU resource. SSDs contain advanced ASICs which provide the flash functions very efficiently on the SSD itself, exploiting internal flash buses and device characteristics, and freeing the server's processor cores for productive use.

A much higher degree of parallelism, balance, and maintainability can be achieved through the use of parallel SSDs over PCIe flash memory subsystems. The flash memory configuration can be adjusted to match workload capacity, bandwidth, and latency requirements with optimized controller/SSD configurations, and servers can be provisioned with higher flash memory capacity using SSDs than with PCIe flash memory subsystems.

It is also much easier to replace an SSD than a PCIe flash memory subsystem—somewhat similar to the difference between installing a memory stick into a USB port on a typical personal computer (PC) vs. opening up a PC to install a graphics card.

The Intel® X25-E Extreme SSD

The leading-edge example of flash memory technology is the 64 GB 2.5" Intel X25-E Extreme SATA SSD. The X25-E SSD has 10 parallel flash channels accessing SLC NAND flash memory, providing both high performance and reliability. With native command queuing to enable up to 32 concurrent operations, each Intel SSD delivers tens of thousands of input/output operations per second (IOPS) at 75 us read latency and write buffering for virtually instantaneous writes.

Each Intel X25-E SSDs has an on-board ASIC which provides write coalescing, wear leveling, and space management, eliminating the need for large consumption of server host processor cycles. Many Intel X25-E SSDs can be integrated and accessed in parallel to provide hundreds of thousands of input/output operations per second (IOPS) per server.



Figure 2. The Intel X25-E Extreme SATA SSD.

Incorporating Flash Memory into Overall System Architecture

A very high degree of parallelism and concurrency control is required in the application and server operating environment in order to effectively utilize the tremendous potential I/O throughput and bandwidth offered by advanced flash memory technology. Also, flash memory driver, controller, device optimization, and tuning are required to match to workload behavior, especially to access size distributions and required persistence semantics.

High-Performance Interconnects

Interconnects have come a long way since Ethernet first became popular in the 1980s. Bandwidth continues to increase while latencies are steadily getting smaller. Today, Gigabit Ethernet is standard on most server motherboards. 10GbE is being used in datacenters mostly as a backbone to consolidate gigabit links and is starting to gain traction as a point-to-point interconnect.

Table 2 lists the useful latencies and bandwidth that can be achieved using both 1Gb and 10Gb Ethernet.

	Latency	Bandwidth
1Gb Ethernet	25.0 us	1 Gb/sec
10Gb Ethernet	6.0 us	10 Gb/sec

Table 2. The latencies and bandwidth of 1 Gb and 10 Gb Ethernet interconnect technologies.

With latencies as low as a single microsecond between server nodes, it is feasible to distribute workloads across multiple servers and to use replication to multiple server nodes to provide high availability and data integrity. Nevertheless, most applications available today were written with the assumptions of high latencies and slow bandwidth. The software to effectively manage data movement at such high speeds while running simultaneously on multiple server nodes is very complex.

Loosely Coupled Scale-Out Architectures

Web and cloud computing build-outs took advantage of commodity x86 servers and GbE, building scale-out applications with tiering, data partitioning, and fault tolerance. These loosely coupled system architectures enabled scaling—but at a high cost in efficiency.

A modern Web 2.0 and cloud data center scale-out system architecture deployment has at the front a Web server tier and an application server tier (which are often merged together), and in the back-end a reliable data tier, often hosted by database servers. Database servers are typically slow and expensive elements of the system architecture. They often operate at very low CPU utilization due to blocking on HDD accesses, lock serialization effects, and low HDD capacity utilization due to having to minimize head movement to reduce access latencies.

Between the Web server tier and the back-end server tier are specialized application services and a content caching tier. The specialized application services tier may perform generic functions such as search, ad serving, photo store/retrieval, authentication, etc., or specific functions for the enterprise. Completing a response to a customer interaction involves accessing a Web server, application servers, database servers, and various other generic and specialized applications and servers.

Datacenters often require that user responses complete in less than $\frac{1}{4}$ second. In order to accomplish this, a DRAM caching tier is often utilized which consists of servers filled with DRAM. The caching servers present a simple (key, value) interface, wherein the client applications provide a key (along with the value associated with the key) for caching, and later retrieve the cached data by supplying the key. Customer information, data retrieved from slow databases, and common user interaction results are often cached in this DRAM tier so they can be accessed very quickly.

Since the performance of a site can often be improved dramatically through extensive caching, many racks of caching servers are often deployed. Each caching server only holds a limited amount of DRAM, so the data must be partitioned among the caching servers. Application designers and administrators need to carefully lay out the data between the caching servers. These caching servers typically operate at very low network and CPU utilization as they are simply storing and retrieving pieces of the relatively small amounts of data they are caching when requested by the various client Web or application servers.

When loosely coupled scale-out architectures are examined closely, it becomes clear that the database and caching tiers suffer from very low utilization, high power consumption, and excessive programmatic and administrative complexity—all contributing to high TCO.

Challenges of Utilizing New Technologies

The new generation of commodity multi-core processors, flash memory, and low latency interconnects offer tremendous potential in Web 2.0 and cloud computing datacenters in terms of inherent performance, power, and space. Unfortunately, when the technologies are just plugged into existing server deployments, the realized benefits are very limited.

The middleware applications driving the database and caching tier lack sufficient thread level parallelism, granular concurrency control, and thread and data affinity management to drive multi-core processors, resulting in very low utilization. Effectively utilizing high performance interconnects requires highly efficient and parallel initiation and completion processing for multi-node server replication, load balancing, and consistency management.

Flash memory offers access times 100x faster than HDD, provides much more capacity with significant lower power consumption than DRAM, and is persistent. Yet it is neither DRAM nor a HDD and has many idiosyncrasies. Not only does flash memory require high thread level parallelism and granular concurrency control to utilize its low access times and high IOPS and bandwidth, it also requires complex space management, caching, persistence control, and wear-level control. In addition, flash memory performance is access-pattern-sensitive, requiring optimization of configurations and tuning of drivers and controllers.

The effort required to effectively utilize these new technologies to solve today's severe performance, power, space, and TCO challenges is significant and the challenges are many. IT teams effectively need to develop highly parallel middleware applications, a high performance operating system, and develop and optimize numerous, specialized configurations.

Adapting or inventing a new deployment architecture which effectively takes advantage of the new technologies is almost a research project, with large development and support costs which are not the core value of Web 2.0 and cloud computing enterprises. Most enterprises neither could nor would want to invest in this level of requisite research, development, integration, and support to realize the lucrative potential of these new technologies.

The Schooner Solution: Tightly Coupled Building Blocks

Schooner Information Technology has pioneered a new generation of data access appliances specifically architected for Web 2.0 and cloud computing datacenters. All Schooner appliances are built on a patent-pending system architecture that incorporates enterprise-class flash memory, multi-core processors, low-latency interconnect, and highly optimized data access and caching applications.

The Schooner architecture was designed from scratch based on extensive workload characterization, system modeling, experimentation, and optimization. The result is highly optimized appliances that are scalable, smart, cost effective, and green.

By leveraging Schooner's higher-level building blocks, datacenter managers are no longer burdened with complex and inefficient integration projects. They can now leverage the advanced technology of Schooner's integrated appliances to quickly and easily manage data growth, decrease business complexity, and cut TCO.

Scalable Hardware Platform

The Schooner hardware platform design, shown in Figure 3, utilizes the most advanced technologies available today, including:

- Intel Nehalem 5550 processors
- IBM System x3650 M2 server platform
- 512GB Flash Memory Subsystem, employing highly parallel flash controllers and Intel X25e SSDs
- High throughput networking: 1/10-GbE for client traffic with Web, application, and specialized servers

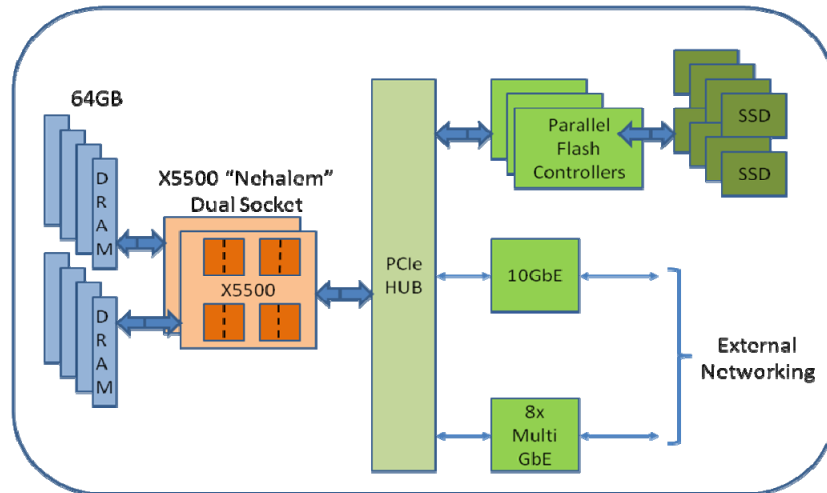


Figure 3. The Schooner hardware platform utilizes the industry's most advanced technologies.

While the architecture and design of Schooner appliances was invented, patented, and developed by the Schooner team, Schooner collaborated closely with IBM throughout the implementation phase. As a result, all Schooner's appliances are built on an IBM-optimized server platform, and are built from the ground up for serviceability. Schooner has also partnered with IBM to provide world-class 24/7/365 single-point-of-contact service and support for all Schooner appliances.

Schooner Operating Environment

The Schooner Operating Environment (SOE) is the central technology for all Schooner appliances. SOE optimizes high-performance multi-core processors, flash memory, DRAM, and low latency interconnect in a highly parallel, low overhead manner to balance system resources and maximize system throughput. The Schooner Operating Environment is depicted in Figure 4.

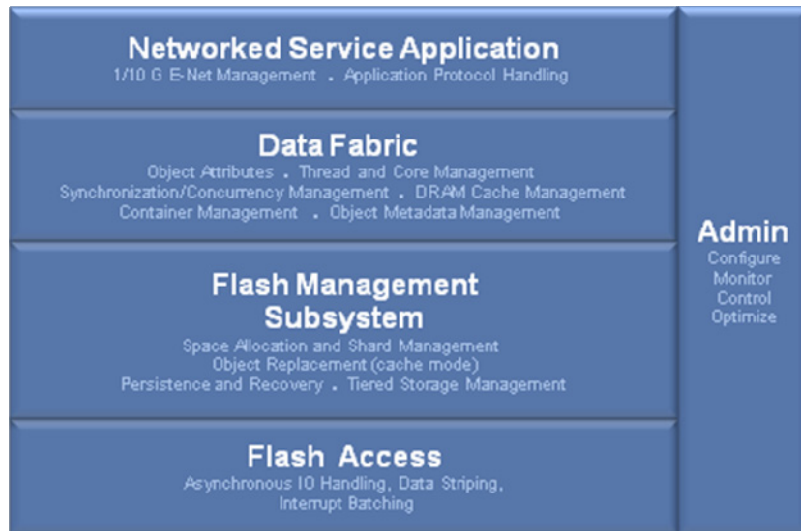


Figure 4. The Schooner Operating Environment is the central technology for all Schooner appliances.

SOE includes the following functionality:

- SOE directly and holistically manages all system resources, including threads, cores, interconnect, and DRAM and flash memory.
- It optimizes key performance metrics, including transactions/sec/watt, transactions/sec/core, and transactions/sec/\$.
- SOE co-exists with Linux on all Schooner appliances.
- The networked data-access middleware application-in-an-appliance compatibly supports the application level protocols, but its implementation is directly integrated with SOE's data fabric.
- SOE's data fabric is optimized to take advantage of Intel Nehalem multi-core architecture, allowing processes to communicate and context-switch without the expense of a kernel call, and with careful scheduling of threads and placement of data for affinity.
- SOE's flash management subsystem manages highly parallel flash memory devices, optimizing data caching, placement, and movement based on workload behavior and characteristic access times and bandwidths.
- SOE's high performance interconnect management, with highly parallel, low-overhead, batched initiation and completion processing, provides optimized client throughput, as well as highly efficient multi-node replication and load balancing.

The SOE is tightly integrated and was developed to use the new generation technologies to their best advantage. This enables each Schooner appliance to make the optimal use of the CPU cores, flash memory, and low latency interconnects to deliver the benefits directly to the datacenter, based on optimized use of resources and maximal application throughput at minimal power and space.

MySQL and Memcached Appliances

Schooner currently offers two product families: the Schooner cache family, with the Schooner Appliance for Memcached; and the database family, with the Schooner Appliance for MySQL Enterprise. Both appliances manage multiple terabytes of data with great speed and energy efficiency, enabling high-performance data access for all Internet-based businesses.



Figure 5. The Schooner Appliance for Memcached

Memcached Challenges

Memcached is a key/value distributed cache commonly used in high-traffic Websites, with many sites employing hundreds of Memcached servers. Originally written by Brad Fitzpatrick in Perl, it was re-written in C and optimized to better serve the increased demands that were being placed on it.

Legacy Memcached installations are typically limited by the amount of DRAM that can be installed in a single server node. If an application requires a total of 512GB of memory to hold the working set of an application, legacy deployments require 16 Memcached nodes at 32GB of DRAM per server node. Since the caching workload must be spread among many legacy Memcached servers, and the legacy Memcached servers have poor multi-core parallelism—the result is low CPU utilization. Most legacy Memcached nodes use 1Gb Ethernet links, which are sufficient because of the limited amount of DRAM and CPU utilization per node.

Schooner Appliance for Memcached

Schooner has developed a new Memcached appliance which is tightly integrated with the Schooner scalable hardware platform and Schooner Operating Environment, providing order-of-magnitude improvements in performance, power, and space. The Schooner Appliance for Memcached is 100% compatible with existing Memcached client applications.

The Schooner Appliance for Memcached incorporates highly optimized Memcached server code implementing the Memcached protocol at the networked application layer. This appliance is tightly integrated with the Schooner Operating Environment, holistically balancing 16-core hyper-threaded Intel Nehalem processing with the network bandwidth of multiple 1Gb and/or 10Gb Ethernet links, and hundreds of thousands of accesses per second of flash memory access capacity through parallel flash memory controllers and SSDs, to and from 512GB of NAND flash memory.

The Schooner Appliance for Memcached achieves very high Memcached throughput and capacity per appliance node. With 512 GB of cached data per appliance, an order-of-magnitude reduction in the number of nodes is achieved, relative to legacy deployments. Furthermore, the higher throughput per Schooner Appliance for Memcached better exploits high bandwidth networking.

Implementing the bulk of the Memcached data set in flash memory has the added benefit of significantly reducing power consumption. By consolidating up to 10 legacy DRAM-based Memcached server nodes, the Schooner Appliance for Memcached reduces power consumption by up to 10x. Performance, network, and power comparison results are shown in Figure 6.




Figure 6. Schooner appliances vs. legacy Memcached: Comparing throughput, network bandwidth, and power efficiency.

Memcached performance is typically measured in transactions per second (TPS) (i.e., the number of get/set requests the Memcached server can service from incoming clients in a second). Figure 6 shows a comparison between the open-source v1.2.6 Memcached running on a 16-core hyper-threaded Intel Nehalem processor and the Schooner Appliance for Memcached with a 16-core hyper-threaded Intel Nehalem processor, 64GB of DRAM, a 10Gb Ethernet adapter, and 512GB of flash memory. The workload uses 4KB objects with a key size of 128 bytes. 95% of the operations are “set” and 5% are “get,” with 5% of the get operations missing in the Schooner DRAM cache.

The performance of legacy Memcached, limited by DRAM capacity, utilizes a single 1Gb network interface, and achieves 10% processor utilization. The Schooner peak throughput is achieved at 100% processor utilization, at which point the 10Gb network interface and flash subsystem are fully balanced. Figure 6 also shows the increase in usable network bandwidth and the improved power efficiency in the Schooner Appliance for Memcached.

Multiple instances of Memcached with different attributes can be configured on a single Schooner Appliance for Memcached. Attributes of a Schooner Memcached instance include size, persistence, store-mode, and replication settings:

- Schooner persistence attributes enable Memcached data to persist across power failures and appliance restarts. The Schooner Appliance for Memcached allows frequently accessed data to be cached in DRAM for maximum performance. Schooner has extended the Memcached protocol

with synchronization commands that allow applications to explicitly control the synchronization of objects cached in DRAM with their copies in persistent flash memory.

- A Schooner Memcached instance can optionally be specified to function as a key/value store, rather than a cache, so that data cannot be evicted unless explicitly deleted.
- Replication in multi-node Schooner Memcached clusters can be optionally specified on a per-instance basis, and the replicas will transparently take over in the event of a failure.

MySQL Enterprise Challenges

With over ten million installations, MySQL is the most popular open-source database. It has become the database of choice for the Web 2.0 market. Transaction processing workloads have accesses to storage that are random in nature and are affected deeply by the number of random I/O operations provided by HDDs (typically only 150-200 random IO per sec per drive). As a result, most MySQL servers must use many 15k RPM HDDs with just 73GB data capacity, and still must place data on the outer regions of the platter in order to minimize head movement and access time. Even with this HDD optimization, the servers are bottlenecked on HDD and processor utilization is very low.

HDDs also exhibit poor reliability due to mechanical moving parts. They represent the largest component responsible for hardware failures. HDDs must therefore be deployed in RAID configurations to eliminate or minimize failure effects, which also reduce the capacity and parallelism available from the available disk slots in a commodity server.

The buffer pool in the database is a portion of main-memory DRAM that is utilized to cache storage blocks. Keeping most or all of the working set of a workload in the buffer pool can help alleviate the slow-down in processing a request by reducing the number of disk accesses. Tuning the data size according to the workload and changing user and application patterns so that the working set fits in the buffer pool effectively is a hard problem to solve. It requires extensive data partitioning between the database servers and optimization of the working set of a database so it resides primarily in main memory (in the database buffer pool) in order to minimize HDD accesses.

Deployment architects must also limit load on a database server due the limited scalability of legacy multi-core processors and the heavy lock and HDD contention. Furthermore, legacy slaves are often unable to keep up or catch up with master shard databases due to limitation of single-thread performance while executing updates from the replication log. This has further impact on overall performance and global consistency of data.

Also, application query mixes must be carefully controlled due to the effect that queries such as table-scans can have on performance due to wiping out the buffer pool. Warm-up and recovery times after planned or unplanned restarts are also painfully long. Simply replacing HDDs with SSDs does not solve the problem due to insufficient parallelism and coarse concurrency control in the operating system. The result is still very low multi-core CPU utilization, low effective transaction throughput, and potential lack of integrity when using SSD write-caches for performance (SSD flash buffers can cause data corruption in the event of a power loss or system restart if not explicitly synchronized).

The Schooner Appliance for MySQL Enterprise

The Schooner Appliance for MySQL Enterprise is tightly integrated with the Schooner Operating Environment and the Schooner hardware platform to fully exploit Intel Nehalem multi-core processors with hyper-threading and 512 GB of parallel NAND flash memory.

The Schooner Appliance for MySQL Enterprise provides consolidation on the order of 8 to 1, with 60% reduction in TCO and large reductions in development, deployment, and administrative complexity. Schooner's administrative tools provide easy configuration, monitoring, and optimization of MySQL deployments. They allow for GUI or CLI initiated installing and configuring of a new instance, setting master-master or master-slave relationships between different instances, doing point-in-time recovery, administering and monitoring a group of database instances, and performance tuning.

Figure 7 provides a comparison between legacy MySQL and the Schooner Appliance for MySQL Enterprise running the DBT2 open-source implementation of the TPC-C benchmark. Legacy MySQL is shown both for a disk-optimized and SSD optimized configuration with dual socket Nehalem quad-core processors.

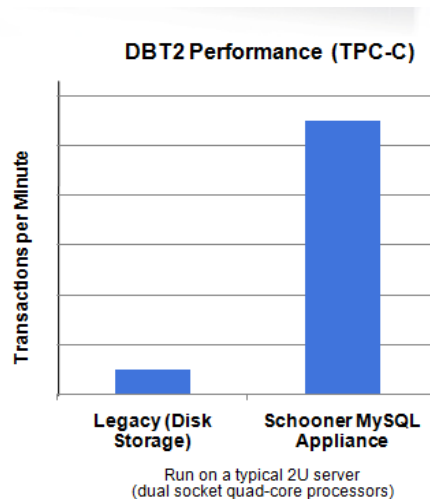


Figure 7. The Schooner Appliance for MySQL Enterprise provides 8x performance increases compared with legacy disks.

Schooner Administrative and Analysis Tools

A powerful set of tools is used to administer, measure, and optimize Schooner appliances. For increased flexibility, customers may choose to administer Schooner appliances using either Schooner-supplied tools, open source, or standards-based tools.

The Schooner Administrator provides a Web GUI for IP assignment, image installation and updates, and event and performance monitoring of the grid and its member Schooner appliances. (The alternate methods of provisioning and monitoring Schooner appliances include CLI, SNMP, Nagios, Ganglia, and Cacti. The Ganglia grid monitoring system provides performance statistics at the grid, group, and appliance level.)

At the Schooner appliance level, Schooner analysis tools instrument and display the following metrics:

- Throughput and response times
- CPU utilization breakdown into networking, flash, and application activity
- Memory usage and miss rates for each level of storage
- Service queue lengths
- Power utilization
- Key and data distributions
- Other general and appliance-specific statistics

These metrics and tools enable configuration and optimization for all customer workloads and deployments.

Benefits of Schooner Appliances

Schooner's appliances provide up to 8x higher throughput than traditional servers. A single Schooner appliance can replace many traditional servers, providing an immediate capex and ongoing opex savings of over 60%.

Both the Schooner Appliance for MySQL Enterprise and the Schooner Appliance for Memcached are 100% compatible with the client applications in use today, allowing for rapid deployment. Because they provide plug-and-play installation, they are easier to initialize and administer than the traditional servers they replace. In addition, the Schooner Administrator employs extensive monitoring and optimization features which are easily integrated with existing management tools.

Schooner appliances employ persistence, replication, and recovery software to deliver enterprise-class reliability and dramatically increase the mean time between failures (MTBF). And finally, all Schooner Appliances are supported by IBM, which provides 24/7/365 single-point-of-contact service.

About Schooner Information Technology

Schooner was founded in 2007. The company is headquartered in Menlo Park, California, with development offices in Hang Zhou, China. Schooner is currently funded by leading technology investment firms Redpoint Ventures and CMEA Capital. To learn more about Schooner, please visit www.schoonerinfotech.com/.



Schooner Information Technology

1350 Willow Road, Suite 101

Menlo Park, CA 94025 USA

Tel: 650-328-4200 Fax: 650-328-4201

info@schoonerinfotech.com

www.schoonerinfotech.com

© 2009 Schooner Information Technology