Objectives

Data storage processing is typically structured as an "I/O path" that takes data from the application to the storage device itself. The I/O path in existing virtualised systems traverses through several layers of the system: the application and middleware layer, the virtual machine (VM) layer, the host operating system (OS) layer and the embedded storage controller firmware layer. In non-virtualised systems, there are no virtual machines and applications are executed on the host OS. This generic layered structure is shown in Figure 1.

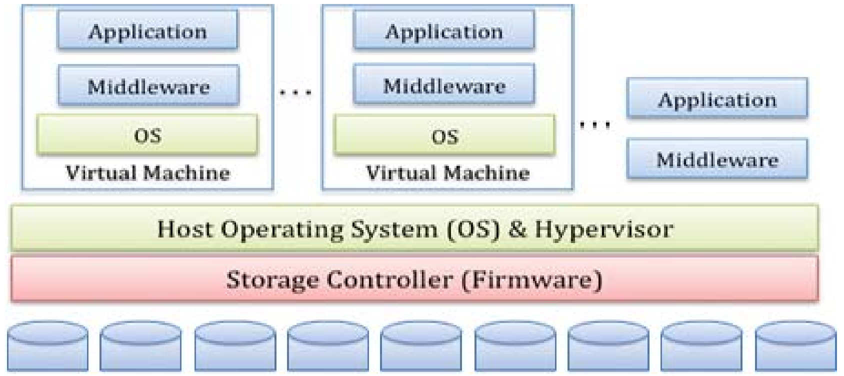

***Figure 1:** Layers in the I/O path of existing systems*

IOLanes aims at addressing issues in modern and future I/O stacks and throughout the I/O path from the application to the storage device. Our approach breaks the I/O stack in four layers: (a) application and middleware, (b) virtual machine, (c) host operating system, and (d) embedded storage controller.

Figure 2 (left) shows a more detailed view of the I/O stack in modern systems. Essentially, IOLanes maps this complex stack on a modern multicore platform (right) dealing with memory, scheduling, protection, and division of functionality issues. This is a central problem for all modern applications that depend on storage, both because of increasing application performance requirements and because of increasing needs in novel functionality as well. Multicore systems with the increasing number of cores per chip have resulted in reduced effective storage bandwidth available per core. The CPU vs. I/O performance gap becomes wider and we need increasingly faster storage to keep CPUs busy with data.

IOLanes studies and addresses the bottlenecks, overheads, and inefficiencies associated with all levels of the I/O stack on multicore architectures. We use the Linux I/O stack as the basis for our work and we will design and build new techniques and methods for addressing the aforementioned issues. This will result in systems that are able to perform multi-GBytes/s I/O in virtualised environments, supporting advanced functionality, and allowing scalability with the number of cores, as multicore architectures evolve over time.

Partners

Foundation for Research & Technology - Hellas (FORTH)
Role: Storage infrastructure, block-level I/O systems, embedded I/O software, storage virtualisation, cluster storage and file systems, operating systems.

Univ. Politecnica Madrid (UPM)
Role: I/O applications, transactional systems, high-performance databases, web and data streaming systems, distributed systems, performance evaluation.

Barcelona Supercomputing Center (BSC)
Role: Storage infrastructure, data placement, heterogeneous storage systems, scalable storage systems, file-system virtualisation, storage caching and prefetching, operating systems.

IBM Israel - Science and Technology LTD (IBM)
Role: Storage controllers, operating systems and hypervisors design and implementation, virtualisation and in particular I/O for virtualised systems, multicore architectures, high-performance networking.

Intel Performance Learning Solutions Ltd. (INTEL)
Role: Distributed systems, infrastructure virtualisation, manageability, service interfaces, integration, Instrumentation, analytics.

Neurocom S.A. (NEUROCOM)
Role: Business intelligence and data warehousing, very large databases, multi-tier architectures, business analysis and requirements gathering.

Contract Number: 248615

BUDGET:	**TIMETABLE:**
Total cost: €4.260.426	Starting date: 2010/01/01
Funding: €2.718.000	Duration: 36 months

This project is part of the portfolio of the **Embedded Systems Unit - G3 Directorate General Information Society**

For more information please check:
http://cordis.europa.eu/ist/embedded

Information Society Technologies

Contact Person: Angelos Bilas
Foundation for Research and Technology - Hellas (FORTH) • Institute of Computer Science • N. Plastira 100 • GR - 700 13
Heraklion, Crete, Greece • Tel.: +30 2810391669 • Fax: +30 2810391601 • Email: bilas@ics.forth.gr • http://www.iolanes.eu/

IOLanes:

Advancing the Scalability and Performance of I/O Subsystems in Multicore Platforms

IOLanes

http://www.iolanes.eu

Scope 1.

Data storage technology today faces many challenges, including performance inefficiencies, inadequate dependability and integrity guarantees, limited scalability, loss of confidentiality, poor resource sharing, and increased ownership and management costs. Given the importance of both direct-attached and networked storage systems for modern applications, it becomes imperative to address these issues. Multicore CPUs offer the promise of dealing with many of the underlying limitations of today's I/O architectures. However, this requires careful consideration of architectural and systems issues and complex interactions in the I/O stack, all the way from the application to the disk.

IOLanes targets three major challenges: (i) dealing with performance and scalability issues of the I/O stack on multicore architectures, (ii) examining on-loading and off-loading tradeoffs on modern multicores for advanced functions that are becoming essential in modern storage systems, such as compression, protection, encryption, error correction, and (iii) addressing I/O performance and dynamic resource management issues in virtualised, single-host environments.

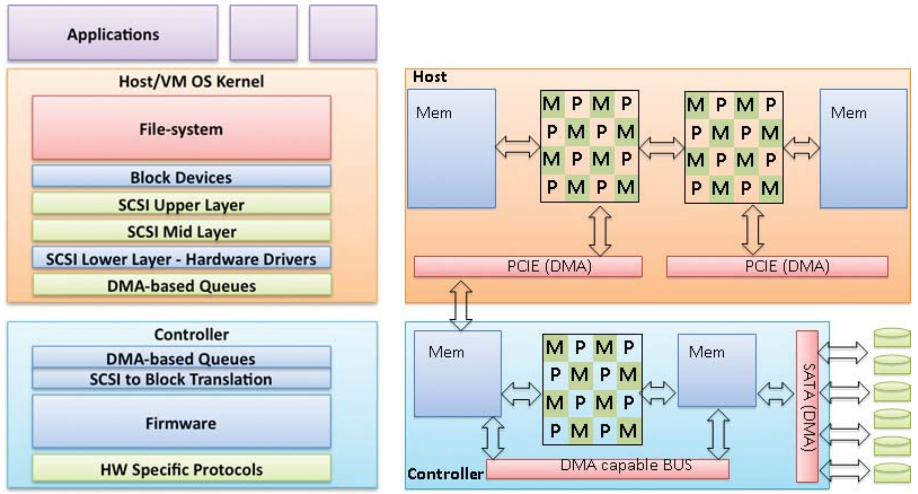

Figure 2: *A more detailed view of the I/O stack in modern systems*

Challenges 3.

- Scaling the I/O stack across all system layers to exploit the CPU performance offered by multicore processors.
- Exploiting the computing power of multicore processors to improve the performance of the storage subsystem via "on/off-loading" mechanisms.
- The interaction of the I/O paths of multiple isolated virtual machines over a single host operating system.

Demand for Storage I/O 4.

Our lives are increasingly becoming dependent on electronic information that is processed and stored in computer systems. Individuals, businesses, and organisations cannot survive in today's competitive economy without the use of digital information stored in electronic data storage systems. Data storage is perhaps the most critical and valuable component of today's computing infrastructures. However, critical and valuable as they may be, existing data storage systems today fall short of applications and users' needs in four main respects:

1. Performance: The storage system continues to be the performance bottleneck in most computer systems, due to the processor-disk performance gap. Today, the recent advent of multicore processors and the steadily increasing number of cores per chip has resulted in a decrease in the effective storage bandwidth available per core. As multicore processors become faster, one needs increasingly faster storage to keep them busy with data. Instead, today, in the case of multicore CPUs, the CPU vs. I/O performance gap becomes wider.

2. Dependability & Data Integrity: For the last three decades there has been a dramatic increase of storage density in magnetic and optical media, which has led to significantly lower cost per capacity unit ($/GB). Furthermore, organisations and individuals have taken advantage of this trend by creating and storing ever-increasing amounts of digital data. Storing and handling these, unprecedented amounts of data, has led to increased failures; and failures in the storage subsystem can be unnerving. According to the US National Archives and Records Administration report, "93% of the companies that lost their data center for 10 days or more due to a disaster, filed for bankruptcy within one year of the disaster". Moreover, according to the same source, "50% of the businesses that found themselves without data management for this same time period filed for bankruptcy immediately". In a recent Storage Networking Industry Association (SNIA) survey, the most important issue for storage administrators was reliability.

3. Data Confidentiality: The sensitivity of digital data for individuals and organisations requires careful protection of all digital assets. Data should always be stored in e.g. an encrypted form to prevent leakage to anyone who can obtain access to the operating systems or physical access to the storage device(s). However, such solutions remain to date exotic, are not used in most systems, and even when they can be used, they are extremely costly in terms of processing cycles and power. Fine-grained access control for on- and off-line data can benefit from architectural support in modern systems.

4. Resource and content consolidation: A main requirement for modern storage systems is to consolidate both storage resources as well as data and content. Consolidation of resources results in higher efficiency and improves costs, whereas consolidation of data and content provides more flexibility to applications and users. However, resource and data consolidation require mechanisms for isolating resources and protecting data that are today only possible by physical separation of resources and data. Future systems should employ extensive virtualization as well as new data protection technologies to ensure that both resources and content are utilized efficiently and properly.

A common theme that underlies these four issues is the overheads incurred in today's I/O stack in modern operating systems and their limitation to scale with the number of CPUs and take advantage of modern multicore CPUs. Thus, a basic enabler for building the future storage systems is to take advantage in the I/O stack of the performance potential of multicore CPUs and at the same time deal with their shortcomings

Impact 5.

IOLanes will have an impact on how storage infrastructures are designed. It will result in improving I/O efficiency and achieving higher I/O rates at lower costs for demanding applications by providing better use of the available resources in modern systems, including storage devices, controllers, and multicore CPUs. Furthermore, this will also lead to better building blocks for scalable storage systems used in modern and future IT infrastructures.