Automated Application Acceleration with F5 Acopia

Leveraging application-specific intelligent file virtualization to optimize data access and maximize performance

Introduction

Fast, reliable access to data is a necessity in enterprise file storage environments today. However, many environments are beset by bottlenecks that slow application performance and negatively impact productivity. While these problems are particularly acute in data-intensive environments such as digital media, electronic design automation, oil and gas, and financial analysis, they are not unique to such environments and impact many diverse applications across a wide range of enterprises.

The Impact of File Storage Bottlenecks

Bottlenecks often arise as a result of contention for a particular storage resource – for example, high-performance computer farms or server grids that process a specific data set or multiple users accessing a popular file. Common storage management tasks can also result in performance challenges on file storage resources.

Bottlenecks lead to application delays, poor response times, and increased latency, resulting in:

- Increased downtime, lost business productivity, and project delays;
- System instability or unplanned application downtime;
- Increased management overhead to troubleshoot and react to delays by reconfiguring existing environments or creating complex new architectures;
- Increased cost in proliferation of redundant files and additional storage in local severs to improve performance.

To date the "conventional" reaction to such challenges has simply been to add additional storage hardware. In many cases this has obscured rather than solved the bottleneck problem, while increasing the capital and operational costs in the process. However, advances in technology have enabled a new approach to this problem.

Extending ILM to Address Storage Bottlenecks

Information Lifecycle Management (ILM) is often thought of, and has most often been implemented, in the context of archiving aged data. While this is certainly a valuable application of the technology, it is a

Table of Contents

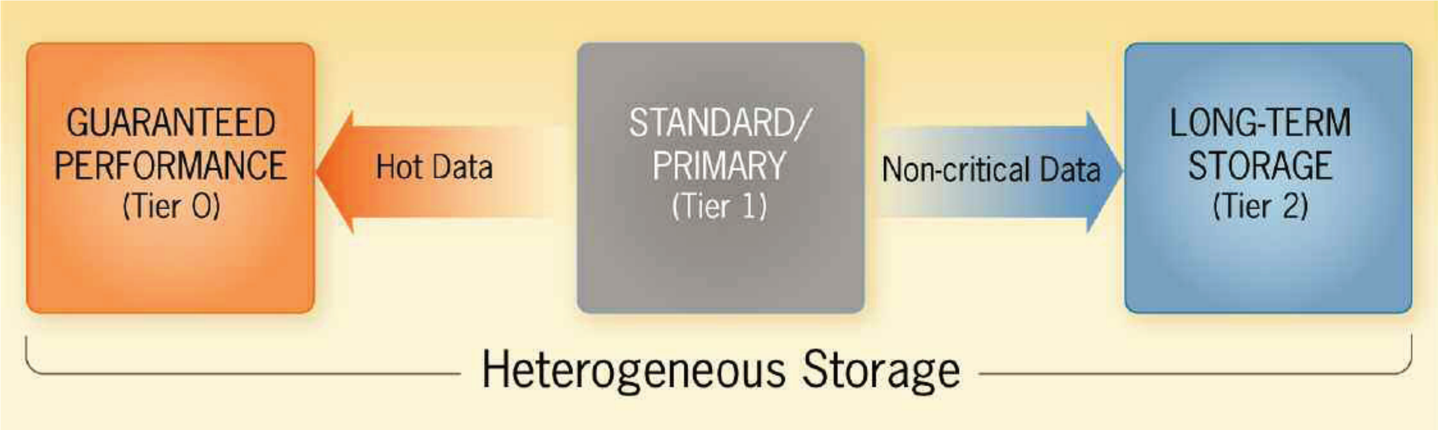

Figure 1*: Three Tier Storage Environment*

limited perspective since ILM essentially involves the matching of business needs with storage capability. The concept of moving data according to its value, age, or performance requirement is commonly accepted and certainly applies to moving nonbusiness-critical data from standard or primary storage (Tier 1) to a cost-optimized archived storage tier (Tier 2). However, the concept applies equally to moving appropriate business-critical data to a high-performance storage tier – Tier 0 – as shown in Figure 1. In fact, F5 Acopia has discovered that there is a powerful positive effect on performance of specific applications when the data associated with them is moved to a high-performance Tier 0. Furthermore, with the advent of inexpensive, commodity-based servers with large, inexpensive internal memory, the price/performance impacts reach extraordinary levels.

Several factors have converged to enable a performance-optimized tier to be implemented both cost-effectively and non-disruptively today:

- Industry standard servers have evolved to the point where they are able to host large memory configurations – 16–128GB;
- Many enterprise environments today have a "hot" data set at any given time that can fit into this 16–128GB memory space;
- Virtualization has evolved to allow the placement of these "hot" data sets onto a high-performance storage tier, in real time, without disrupting access and without requiring application or file system reconfiguration.

Performance-Optimized Storage Tier

By augmenting a classical two-tiered storage architecture with a third high-performance storage tier comprised of standard server hardware, and by using real-time virtualization technology to dynamically place appropriate data sets on this storage tier, enterprises can realize significant performance gains at a fraction of the cost of alternate approaches.

Consider the example of a typical development environment as shown in Figure 2a. A NAS (Network Attached Storage) device hosts a variety of data; some of it is accessed by general applications with a subset of the data accessed by grid applications where performance is at a premium. During normal operations these different applications contend for the storage resource, creating a bottleneck. This bottleneck becomes more pronounced during times of peak processing of the grid applications, and performance is severely degraded for both application types.

Now consider the environment shown in Figure 2b. The NAS storage tier has been augmented by a high-speed performance tier comprised of standard server hardware. The heterogeneous file storage environment is virtualized by an intelligent switch capable of enforcing real-time storage tiering policies. Data that is accessed by the general applications remains on the NAS device as before. However, now data that is processed by the grid applications is dynamically placed on Tier 0. Contention is removed and the bottleneck is alleviated. The grid applications

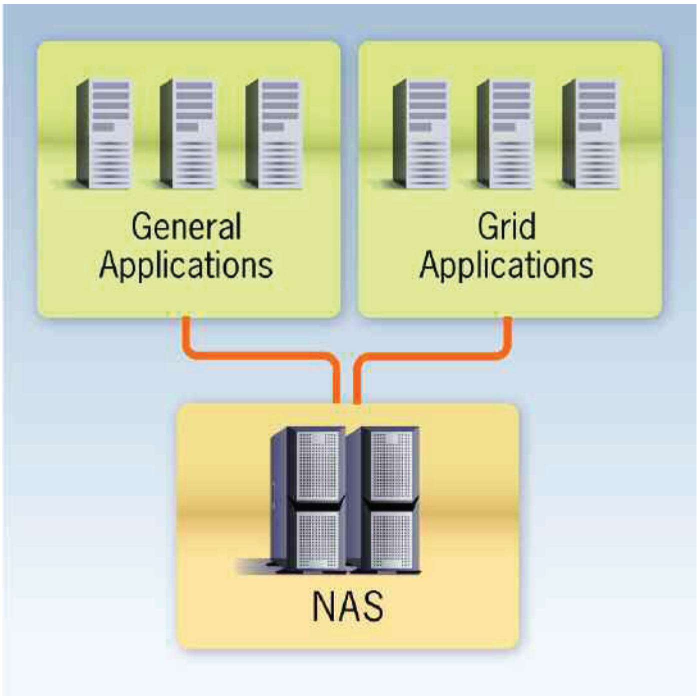

***Figure 2A**: Typical Environment*

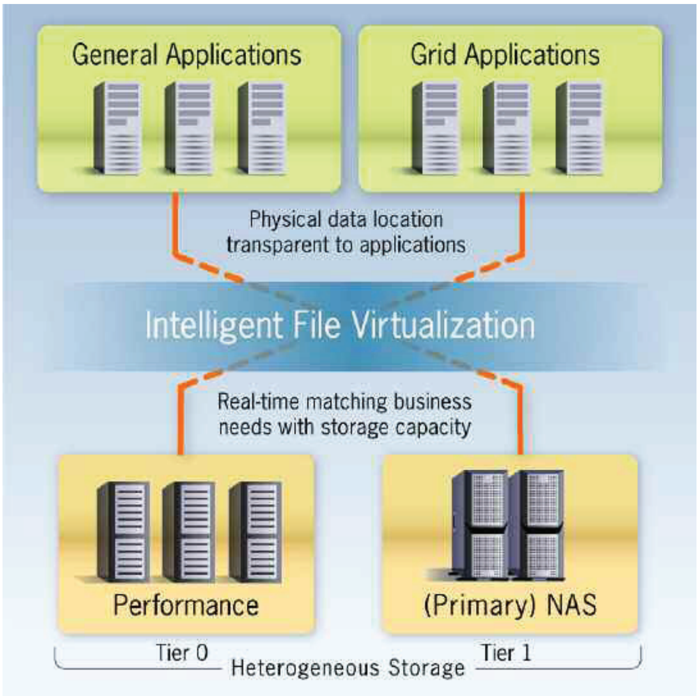

***Figure 2b**: Tier 0 Environment*

are accelerated as they are now served out of a dedicated, high-performance tier with reduced latency and improved throughput. Performance of the general applications is also ensured as they no longer contend for the same resources as the grid applications.

F5 Acopia has conducted internal tests to investigate the effects of Tier 0 on application performance. A test environment was created to match the configuration shown in Figure 2b. Primary storage consisted of a popular mid-range NAS system using high-performance fiber channel disks. The cost of this storage tier configuration was approximately $100K. The performance-optimized Tier 0 consisted of a single, popular off-the-shelf server configured with 32GB memory and dual Intel® processors. The cost of this tier was under $10K.

A simulated grid application was run across 50 application servers using the NFS protocol. This workload accessed a 20GB file set. Two tests were run. In the first test, the grid application file set was served by the primary NAS tier. In the second test, the high-performance tier served the same file set.

The performance results are shown in Figure 3. When the grid application file set was served from the high-performance tier, the grid application throughput increased by almost 30 times compared to using the primary NAS tier. Given the cost difference between the primary and performance storage tiers, the performance tier was 300 times more cost-effective than the primary NAS tier storage when serving this workload.

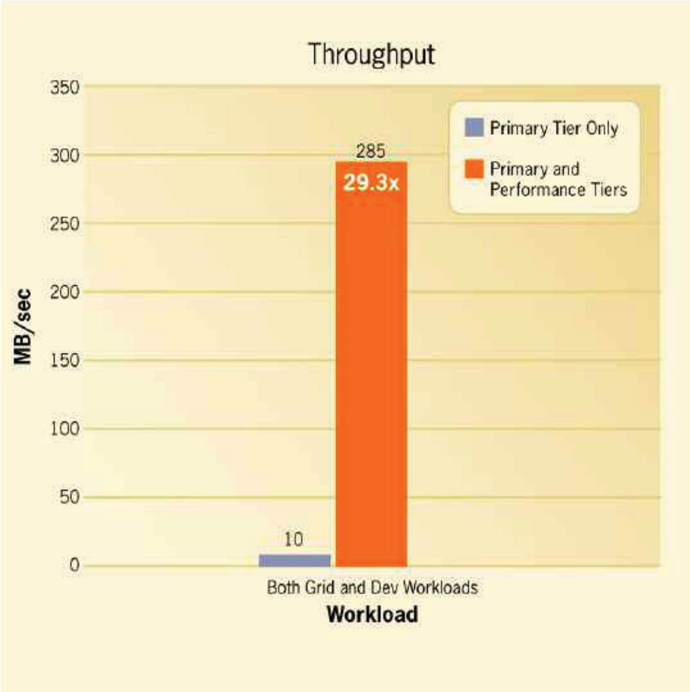

***Figure 3**: Performance Test Results*

Besides the performance and cost benefits of this approach, it is also far less complex and less disruptive than conventional approaches. Additional storage tiers can be added without requiring administrators to reconfigure applications and storage to accommodate new storage devices. More importantly, no service outage is associated with these changes, ensuring productivity is sustained and project schedules are met.

Note that a third tier could be added to the environment shown in Figure 2 that would store data cost-effectively once it had moved beyond a stage of active use, e.g., upon completion of a development project. Figure 4 illustrates how the real-time virtualization engine actively manages this lifecycle over time, using the most appropriate storage tier for different parts of the data life-cycle for maximum efficiency and cost-effectiveness.

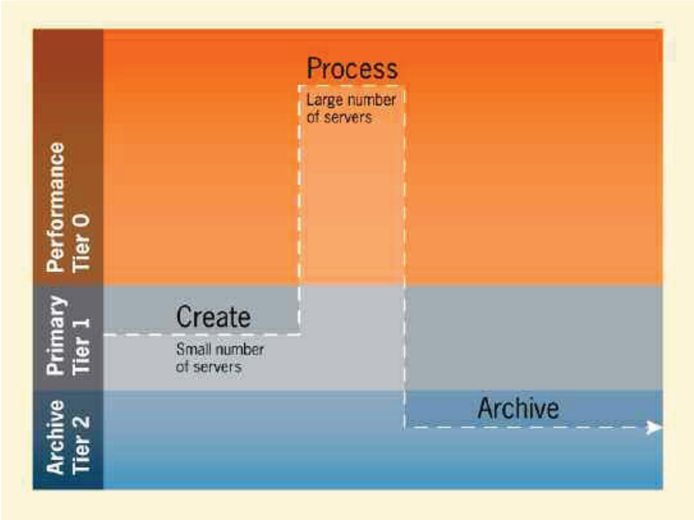

Figure 4: *Matching Application Needs in Real Time*

Summary

Bottlenecks in today's storage environments impact productivity and service levels across many different enterprises. The emergence of intelligent file virtualization, together with advances in standard server technology, allows organizations to address these bottlenecks efficiently and cost-effectively. Now organizations can improve productivity, reduce cycle times, and process more tasks with minimal disruption to their existing environments.

F5 Networks, Inc.
41 Wellman Street
Lowell, MA 01851
(978) 513-2900 Voice
(978) 513-2990 Fax
www.f5.com
info@f5.com

F5 Networks Asia-Pacific
+65-6533-6103 Voice
+65-6533-6106 Fax
info.asia@f5.com

F5 Networks Ltd. Europe/Middle-East/Africa
+44 (0) 1932 582 000 Voice
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

F5 Networks Japan K.K.
+81-3-5114-3200 Voice
+81-3-5114-3201 Fax
info@f5networks.co.jp

Part No. DS-Acopia AAA 0907