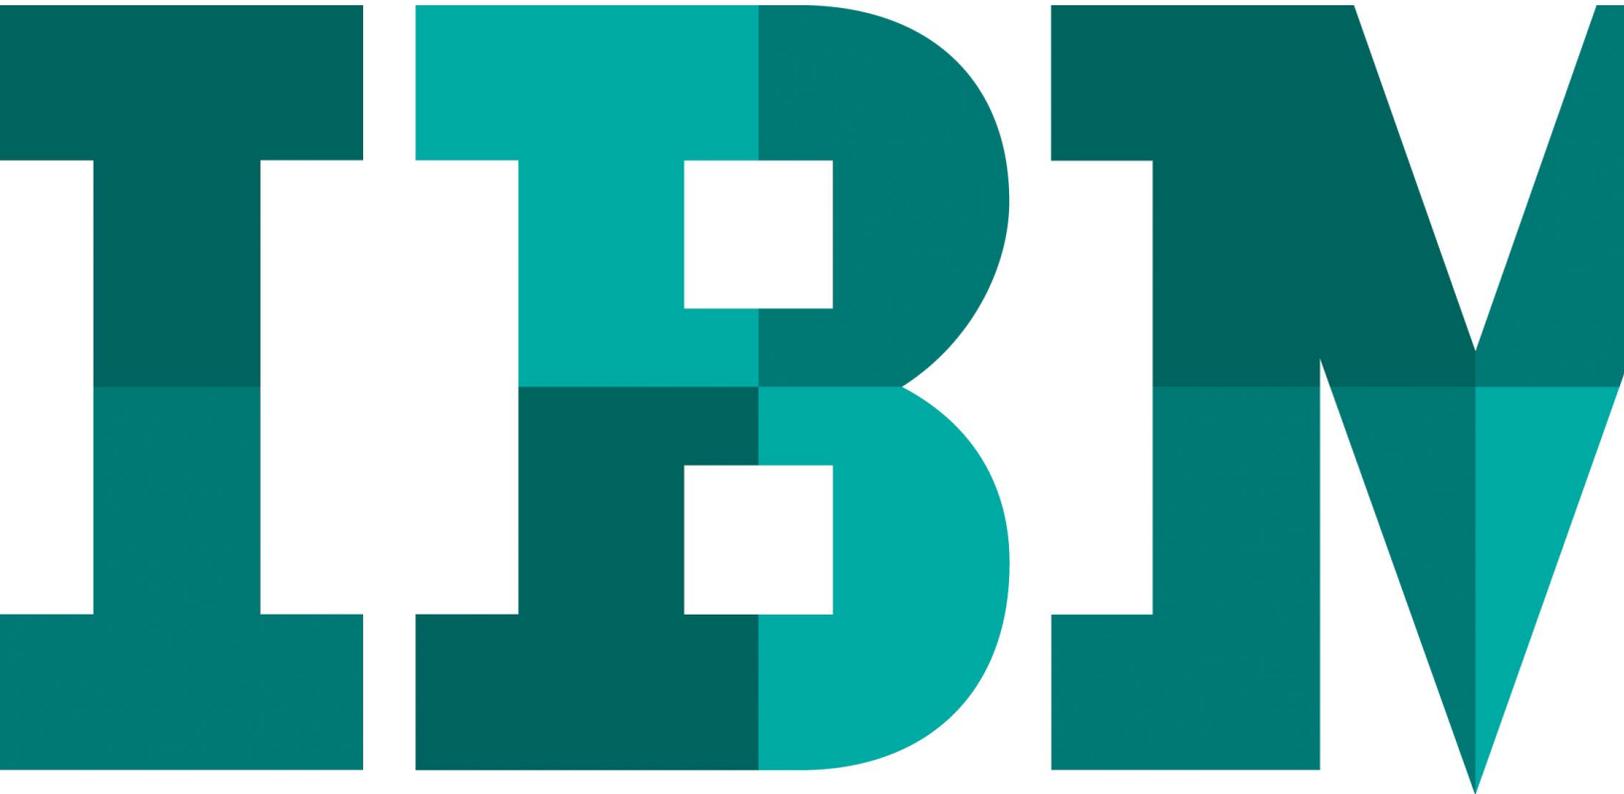


# IBM storage and Actifio copy data virtualization

*Redefining DevOps by improving the accessibility, control, and cost of putting the right data in front of the right people*



## Behind schedule and over budget

The pressure to deliver new applications faster is unrelenting. In addition, 72 percent of software development projects run behind schedule and 54 percent run over budget, according to Panorama Consulting.<sup>1</sup> IBM and Actifio have been working together to solve these challenges. The joint, pre-validated enterprise development and operations (DevOps) solution that resulted from this collaboration is based on IBM storage and Actifio copy data virtualization software, and can improve the accessibility, control and cost of provisioning application development projects. The solution combines strengths in enterprise class data storage from IBM with best-of-breed copy data virtualization technology from Actifio, making management and movement of application data across flash, disk or optional cloud-based media quick, easy and—most importantly—low-cost.

## Application development challenges

Think about the common scenario of provisioning an instance of a production database for a development team. Depending on the size of the database, the process overhead and the availability of storage, that effort could take days or, in the extreme, weeks. This time lag from the point of request to the point of provisioning creates a situation known as data “drift” that prevents developers from using current data and schemas in their work—with the potential for substandard results. Considering that the average enterprise database administrator manages approximately 45 databases, the problem is magnified in an enterprise environment.

Then consider the number of process steps, including approvals, required to facilitate what at first may seem to be a straightforward task. Days can quickly turn into weeks, particularly as work crosses organizational boundaries—from development to database to infrastructure and back. The inefficiencies that result can be significant and compounded based on the number of applications and development and test teams.

Project delays can lead to delays in business acceptance and translate into missed launch deadlines. Those delays greatly impact customer satisfaction. In addition, customer satisfaction can suffer from bugs introduced by lags in the overall process or by shortcuts such as testing against dummy data. Once an application is released into production, unfortunately, these delays are not eliminated. In fact, the change management process will experience a very similar set of delays as each change request triggers legacy approval and provisioning processes.

Infrastructure costs can easily spiral out of control in these environments, particularly as database sizes continue to grow. Consider an organization with three development teams and three test teams. Starting with a 3 TB database in production, the teams must provision a minimum of 18 TB of additional storage to support the application development and test instances of the database. That’s a midsized storage array which would cost a very large sum along with the added costs of the licensing, maintenance fees, and professional services required to set up and deploy the storage array. Additional costs would come from the power draw, Fibre Channel port licenses, operational expenses and other miscellaneous costs of infrastructure. Add a second 3 TB database into the equation and costs become the equivalent of a second array. For many organizations this problem could involve dozens or hundreds of applications.

## The DevOps Methodology

DevOps is a software development methodology emphasizing collaboration between developers and infrastructure teams. DevOps strives for rapid application development and continuous integration through industry-accepted development methodologies, process management and automation. However, organizations still face the challenge of improving automation to efficiently support the full development lifecycle.

Over-provisioned capacity, slow data access and manual, time-intensive processes are still common in many enterprise data environments. Moreover, multiple constituencies across an organization have differing goals and responsibilities:

- Development and testing teams care about fast access, self-service, point-in-time bookmarking, data protection and meeting development deadlines.
- Security teams care about data sprawl, compliance with data privacy mandates, access control and masking of data copies, and audit trails.
- Operations and infrastructure administrators care about storage and compute resource allocation, role-based access control and management, and multi-hypervisor and cloud usage.
- The finance organization cares about controlling costs through storage consolidation, efficient storage utilization and reduced energy cost.

Businesses that fail to fully implement DevOps initiatives risk falling behind existing competitors and new entrants. Organizations that are slow to act can experience impacts on their time to market, revenue growth, customer satisfaction and customer retention. Enterprises are looking for new approaches and technology to help them achieve their goals and keep up with the changing needs of their business.

### DevOps solution from IBM and Actifio

The DevOps solution from IBM and Actifio improves the accessibility, control and cost of putting the right data, in the right format, in front of the right people. The solution combines the strengths of IBM enterprise storage with advanced copy data virtualization technology and automation from Actifio. The joint solution provisions virtual copies of application data instead of creating separate physical copies, dramatically reducing storage consumption without compromising performance or control.

The Actifio platform uses built-in features of IBM® System Storage® solutions to address the data management, availability, and protection of test data typically handled by point tools for functions such as backup, snapshot, clones, replication and de-duplication. By virtualizing copy data—and creating a single “gold copy” available for instant access and recovery—the Actifio solution can dramatically improve service levels while reducing the operational overhead and costs associated with separate storage management tools.

The combination of patented Virtual Data Pipeline (VDP) technology from Actifio with IBM System Storage solutions allows you to create virtual copies of your data with nearly instantaneous access and minimal additional storage requirements per copy. Flash technology, in particular, has traditionally been too expensive to deploy for development and test environments, particularly for very large workloads. Now, with the IBM-Actifio DevOps solution, enterprises can use service level agreements (SLAs) to manage the development lifecycle of their applications and seamlessly match performance requirements to lifecycle stage(s). The DevOps solution can support large numbers of concurrent developers, while being more cost-effective than disk.

The DevOps solution from IBM and Actifio allows all development and quality assurance teams to access virtual copies of production data almost instantly. More importantly, a central administrator for the Actifio solution can specify role-based access control for each virtual copy. Additionally, the administrator can specify the data masking operation that should be used to protect sensitive data before it can be accessed by developers. Those copies can be accessed instantly using a self-service or automated update process.

### Actifio copy data virtualization

Actifio delivers copy data virtualization to more than a thousand global enterprise customers and service provider partners in more than 30 countries. Actifio helps enterprises keep development data current, provision rapidly, automate complex

workflows such as data masking, enable self-service, and do so in a storage capacity-efficient manner. The Actifio approach begins by virtualizing the data in question. Using Actifio's copy data management platform, the process incrementally captures production data on a predefined schedule and maintains it in a rapid-access snapshot pool as well as a long-term, space-efficient pool. In addition to providing nearly instant access to any point-in-time copy of that production data for resiliency purposes, the Actifio solution offers the ability to provision additional virtual, read-write copies for development and test purposes. With Actifio you can:

- Provision read/write virtual copies on demand
- Automate application development workflows and data masking
- Streamline change management with role-based access
- Provision read/write copies of any size database in minutes with minimal storage usage

## IBM Storage

IBM FlashSystem all-flash storage and the Storwize family of storage solutions together provide a high performance, yet cost effective, foundation for the IBM and Actifio DevOps solution.

### IBM FlashSystem

IBM FlashSystem® is a family of all-flash storage arrays engineered to address the most demanding enterprise performance, reliability, comprehensive feature set and cost requirements. IBM FlashSystem 900, the workhorse building block family member utilized in the DevOps solution, is designed to provide extreme performance, low latency and data protection for business-critical application acceleration.

Powered by IBM FlashCore™ technology, IBM FlashSystem 900 delivers the extreme performance, enterprise reliability and operational efficiencies required to provide enterprise-grade targeted application acceleration. For storage environments already equipped with robust storage management capabilities—such as those provided by IBM Spectrum Virtualize™ or

IBM Spectrum Scale™—the ultra-low latency, high performance and ease of implementation offered by IBM FlashSystem 900 make it an ideal choice to transform a business-critical DevOps process into an engine of business growth and competitive advantage.

IBM FlashSystem 900 is composed of up to 12 IBM MicroLatency® modules that provide high-density, high-resiliency storage capacity. IBM FlashSystem 900 can scale usable capacity from as low as 2 TB to as much as 57 TB in a single system. The MicroLatency modules also support an off-load AES-XTS 256 encryption engine and full hot-swap and storage capacity scale-out capabilities. The system uses enterprise-class, two-dimensional flash RAID technology, leveraging both IBM Variable Stripe RAID™ and system-level RAID 5. Variable Stripe RAID maintains system performance and capacity in the event of partial or full flash chip failures, helping reduce downtime and forestall system repairs. System-wide RAID 5 also helps prevent data loss and improves availability.

### IBM Storwize

For some enterprises, especially small and midsized organizations, the initial cost of solution acquisition is a key hurdle. These enterprises can turn to the IBM Storwize® family. With Storwize solutions, IT architects can maximize cost flexibility, especially up-front deployment expenses, while gaining enterprise-grade storage services, including one of the most powerful, mature and comprehensive storage virtualization solutions on the market. At the same time, they can gain the high performance of hybrid and even all-flash configurations.

Built with IBM Spectrum Virtualize software, IBM Storwize V5000 delivers outstanding value, performance, availability, reliability and virtualization capabilities to the DevOps solution configuration. It also includes built-in functions such as IBM Easy Tier® technology, which optimizes flash and hard disk drives to deliver extraordinary levels of efficiency and high performance. Storwize V5000 delivers sophisticated capabilities that are easy to deploy while helping control costs.

## Solution reference architecture

The DevOps solution architecture jointly developed by Actifio and IBM is designed to:

- Maximize throughput without requiring significant changes to any of the underlying components or processes
- Provide a tiered storage model that allows for high-performance storage with IBM FlashSystem all-flash arrays and a capacity-optimized tier using Storwize V5000 storage
- Improve accessibility and performance
- Simplify control
- Lower the overall cost of storage deployments

The solution stack is deployed, nondisruptively, side-by-side with the application development and test environment. Actifio processes automatically discover applications and hosts and ingest them through the IP network. After the first full capture, Actifio technology continues to ingest changes incrementally by change block tracking at the block level, storing in the application's native format.

A typical development/test environment connected to the DevOps solution from IBM and Actifio will consist of:

- One or more subnets for development and test systems
- Physical servers
- Virtualized server stacks

Copy data ingest—frequency, type, retention time, replication and more—is governed by administrator-defined SLAs. Once the data is captured, SLAs also define the location for the storage of that data. A minimum of two pools of storage are defined:

1. A rapid snapshot pool, created on IBM FlashSystem storage, delivers instant access to any retained point-in-time of data.
2. A capacity-optimized pool, created on Storwize storage, is used for longer term and less demanding storage.

The power of the solution lies in its ability to manage point-in-time copies across these pools. This tiered approach provides optimal use of flash and disk per unit cost. Organizations with development or test groups in multiple geographical locations can all use a single instance of the DevOps solution if they have sufficient bandwidth. This allows for the centralization and management of distributed development teams' data to reduce total storage requirements and enhance resiliency.

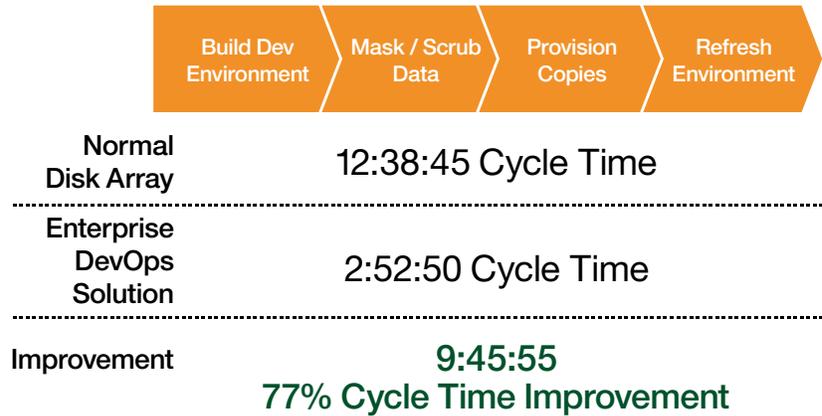
## Delivering real business benefits

The DevOps solution from IBM and Actifio delivers the benefits of rapid test data provisioning and management, capacity efficiency and high-performance/low-latency storage. Other benefits include:

- Rapid access to initial copies from production
- Read/write copies available instantly to authorized users
- 20 times performance increase over conventional storage solutions for all operations<sup>2</sup>
- Rapid refresh based on updates to production
- Faster time to market by removing provisioning bottlenecks
- Increased potential number of development projects per year
- An infrastructure footprint that is six to ten times smaller than conventional storage solutions required for DevOps, based on Actifio experience

The application data lifecycle is dramatically improved with the IBM-Actifio DevOps solution<sup>2</sup>:

- Up to six times improvement in data ingest time compared to standard disk arrays
- Twenty times faster performance in standard development and testing random data access patterns
- Less than one minute mounts of development copies
- At least three times improvement in I/O-intensive operations such as data masking



The following table summarizes many of the benefits provided by the DevOps solution from IBM and Actifio:

<p><b>Development</b></p> <ul style="list-style-type: none"> <li>• Near-instant virtual copies</li> <li>• Self-service features</li> <li>• Automated online virtual database</li> <li>• Test data protection</li> <li>• Point-in-time bookmarks</li> </ul>	<p><b>Security and data control</b></p> <ul style="list-style-type: none"> <li>• Masked data</li> <li>• More control over who can access copies</li> <li>• Increased audit traceability</li> <li>• ITCS 104 validation</li> </ul>
<p><b>Operations</b></p> <ul style="list-style-type: none"> <li>• Role-based access control</li> <li>• Storage efficiency</li> <li>• Multi-hypervisor, cloud</li> <li>• Representational state transfer (RESTful) APIs</li> </ul>	<p><b>Finance</b></p> <ul style="list-style-type: none"> <li>• Consolidated storage footprint</li> <li>• Efficient storage utilization</li> <li>• Reduced energy costs</li> <li>• Reduced database administration, development provisioning and wait time</li> </ul>
<p><b>Copy data services</b></p> <ul style="list-style-type: none"> <li>• Rapid access to initial copy from production</li> <li>• Read/write copies available instantly to authorized users</li> <li>• Bookmarks used for rapid reset</li> <li>• Rapid refresh based on updates to production</li> </ul>	<p><b>Storage throughput, operations and logical unit number (LUN) utilization improvement</b></p> <ul style="list-style-type: none"> <li>• 20 times performance increase for all operations</li> <li>• Faster time to market by removing provisioning bottleneck</li> <li>• Increased potential number of development projects</li> <li>• 6 – 10 times smaller footprint required for DevOps, based on Actifio experience</li> </ul>

## DevOps solution validation

Actifio conducted extensive validation testing on the joint IBM-Actifio DevOps solution to confirm that IBM FlashSystem performance was uncompromised during Actifio copy data virtualization processes, and, in fact, FlashSystem provided 20 times better performance than conventional storage.<sup>2</sup>

### Validation test environment

IBM and Actifio developed the DevOps solution reference architecture and then ran tests using a test bed equivalent to standard application development environments to validate the interoperability and performance of the solution components. The test configuration consisted of two test servers connected to a DevOps solution instance using a Fibre Channel switch. The two test servers were configured identically with the following hardware and software:

- Lenovo x3550 M3 server
- Dual Intel Xeon 2.60 GHz 6-core processors
- 160 GB RAM, 2 x 2.5-inch 300 GB, 15k SAS drives
- QLogic Dual Port 8 Gb Fiber Channel adaptor card
- VMware ESX v5.x

The DevOps solution stack itself consisted of three main components:

- Actifio CDS Appliance
- IBM FlashSystem 900 all-flash array
- IBM Storwize V5000 storage

### Validation test results

A two ESX server test environment was able to drive the equivalent of 31 concurrent, self-contained Microsoft SQL Server systems—each instance with its own virtualized 250 GB database running an online transaction processing-type benchmark and pushing a total of 400,000 input/output operations per second (IOPS) with sustained 2.5 millisecond latency when using IBM FlashSystem storage. More instances would increase the number of currently running test servers.

In the DevOps process, creating the first copy of the primary data source on the DevOps system is a one-time, up-front operation. The DevOps system holds the snapshot only long enough to determine changed blocks, so snapshot time is significantly reduced. When IBM FlashSystem storage powered the validation test environment, ingest time was improved by up to six times compared to disk-based storage. When a standard development and test environment is used with a 65:35 read:write ratio and fairly random data access patterns, the IBM FlashSystem-based DevOps solution configuration provided nearly 20 times faster overall DevOps performance than disk storage.

## Accelerating application development

As the number of new applications being developed continues to increase, enterprises struggle with quality assurance, release management and infrastructure support. Reducing the time to provision instances of production applications for development, test and quality assurance, along with decreasing the lag between live production and those test instances can help address these challenges.

IBM and Actifio have collaborated to develop a DevOps data management and provisioning solution based on copy data virtualization with IBM FlashSystem and Storwize storage systems. The DevOps solution from IBM and Actifio moves data management out of the data path, optimizing the performance and cost benefits that can be achieved by deploying IBM storage. The solution is designed to transform the way organizations manage their data and accelerate application development cycles. The solution offers a relatively simple, cost-effective implementation path for enterprises currently struggling with copy data management issues. And for those who have already implemented Actifio technology, adding IBM System Storage solutions to the mix accelerates the many benefits you have already come to expect.

## For more information

To learn more about the DevOps solution from IBM and Actifio, please contact your IBM representative or IBM Business Partner, or visit the following website: [ibm.com/systems/storage/flash](http://ibm.com/systems/storage/flash)

Actifio virtualizes data, enabling businesses to protect, access and move their data faster, more efficiently and more simply. Learn more about the enterprise-class Actifio copy data virtualization platform at [www.actifio.com](http://www.actifio.com)

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. For credit-qualified clients we can customize an IT financing solution to suit your business requirements, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: [ibm.com/financing](http://ibm.com/financing)



---

© Copyright IBM Corporation 2015

IBM Systems  
Route 100  
Somers, NY 10589

Produced in the United States of America  
October 2015

IBM, the IBM logo, ibm.com, System Storage, Easy Tier, IBM FlashSystem, IBM FlashCore, IBM Spectrum Scale, IBM Spectrum Virtualize, MicroLatency, Storwize, and Variable Stripe RAID are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided. Virtual Data Pipeline (VDP) technology is not an IBM product or offering. Virtual Data Pipeline (VDP) technology is sold or licensed, as the case may be, to users under Actifio’s terms and conditions, which are provided with the product or offering. Availability, and any and all warranties, services and support for Virtual Data Pipeline (VDP) technology is the direct responsibility of, and is provided directly to users by, Actifio.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

<sup>1</sup> Panorama Consulting Solutions, *2014 ERP Report, 2014*  
<http://panorama-consulting.com/resource-center/2014-erp-report/>

<sup>2</sup> See the “validation test results” section of this document for details on these results



Please Recycle