CATALOGIC SOFTWARE

WHITE PAPER

Copy Data Management for a Hybrid Cloud

Technology validation shows how customers can drive dramatic CapEx and OpEx savings by leveraging hybrid cloud for automated disaster recovery

Contents

Executive Summary

What if you could super charge your organization's existing infrastructure? What if you could automate critical business operations such as disaster recovery, test/dev, DevOps and reporting/analytics? What if you could accelerate your competitive advantage while cutting both CapEx and OpEx? All of this is available to you by taking advantage of IBM Cloud's SoftLayer infrastructure and Catalogic ECX copy data management software.

Speeding time to market, increasing reliability and simplifying operations are long standing IT goals, and the cloud promises to deliver on all of them. However, the challenge has been how to automate and orchestrate existing business processes to take advantage of nearly limitless cloud-scale compute power without overloading staff or busting budgets.

Organizations now have a clear path to these goals by deploying Catalogic ECX software in conjunction with IBM Spectrum Virtualize™ software, the IBM Storwize® family of storage systems and the SoftLayer cloud.

The technology validation effort presented in this paper shows how an organization can save more than three times the hard costs and hundreds of hours of IT time automating their disaster recovery (DR) processes to take advantage of the hybrid cloud. Further, by running the DR testing daily, the user can be assured of a far more robust DR strategy. The surprising simplicity of this effort shows how the same operational model can apply for additional use cases such as DevOps and business reporting/analytics.

The result: easier data access, improved data availability, increased competitive advantage and lower CapEx and OpEx.

Introduction

Industry trends show that a vast data center transformation towards a hybrid cloud model is inevitable due to the overwhelming promise of reduced costs and increased business agility. While IT organizations with traditional data center operations will be reluctant to move mission-critical production environments to the cloud too rapidly, business operations that rely on copies of production data are ideal for a hybrid cloud model.

Many business operations require access to copies of production data, ideally on a daily basis. In a traditional data center environment, managing how these copies are supplied to various stakeholders is an expensive and time-consuming job, and IT constantly struggles to keep up with the data access needs of the organization. The result is that data can take days or even weeks to provision and is often stale by the time it gets to users. This means reports are not up to date, critical insights can be missed, and organizational agility is impeded. In a worst-case scenario, vital data can be lost because a snapshot or replication process failed or was overlooked.

What is IBM Softlayer?

SoftLayer, an IBM Company, provides cloud infrastructure as a service from a growing number of data centers and network points of presence around the world. Customers range from Web startups to global enterprises. Products and services include bare metal and virtual servers, networking, turnkey big data solutions, private cloud solutions, and more.

The acute pain of today's copy data management challenge and the great promise of the hybrid cloud create the "perfect storm" to justify modifying data center operations to incorporate hybrid cloud environments in order to drive a more agile organization. Before moving operations that use copy data — such as disaster recovery — to the cloud, IT organizations first need to address the challenges that exist in their on-premises environments.

The Challenge AND the Solution

Traditionally, finding and then using data copies for specific business use cases has been a labor-intensive process that spans multiple storage silos represented by a variety of system administrators. Scripts are typically written to move the data around the enterprise in an effort to automate the process, but the management of scripts has proven to be time-consuming and error prone.

By demonstrating the benefits derived in a single use-case – Automated Disaster Recovery testing – this paper demonstrates how the combination of technologies from IBM, SoftLayer and Catalogic can remove the operational barriers facing IT organizations so that they can easily reap the benefits of a hybrid cloud infrastructure.

What are ECX Copy and Use Policies?

Catalogic ECX defines operations as Copy and Use policies. A Copy policy manages the creation and movement of data copies using IBM FlashCopy snapshots and Global Mirror replication. Copy policies determine location, frequency and retention of copies.

Use policies manage the use of data copies, including mapping snaps and replicas to a target server or recovering volumes in place.

Solution Overview

The example production environment configured for testing included online production applications and data sets that were created using live application data. By simulating an actual user environment, the test results can be extrapolated to demonstrate the value of the solution in a range of real IT environments and use cases.

The simulated production environment was housed in a test lab facility located in New Jersey, where a number of virtual machines were running connected to a Storwize V7000 system. The complete environment consisted of over 100 virtual machines connected to 20 TB of production storage. Testing ran over the course of 60 days simply to make sure that the value stood the test of time. It could have been completed in a day.

Copy and Use Data policies for the VMs were created within Catalogic ECX. Daily replication of the VM data was performed using IBM Global Mirror between the New Jersey based Storwize V7000 and a target array located in the San Jose data center of Equinix, a colocation facility that allows customers to deploy their own infrastructure within its data center.

The neighboring SoftLayer data center served as the location for the VMware ECX hybrid cloud environment that was used to instantiate the replicated VMs to spin up and test the DR environment.

The detailed setup was as follows:

Simulated Customer Production Environment (New Jersey)

- ESX 5.5 – 4 core & 64GB of RAM
- Storwize V7000 array
- Catalogic ECX, deployed as a virtual appliance
- Application VMs hosted on Storwize
  - Exchange 2010
  - Oracle Database server
  - SQL database server
  - Active Directory Server

A diagram of the implementation is shown in Figure 1.

In addition to establishing policies for automated replication on a per-VM basis, the testing also included the automated use of the replicated VM data in a simulated DR test scenario. Testing included full validation of the ability to instantiate the production applications within the SoftLayer/Equinix environment.

Figure 1: Test Solution Implementation – Simulated Customer Deployment

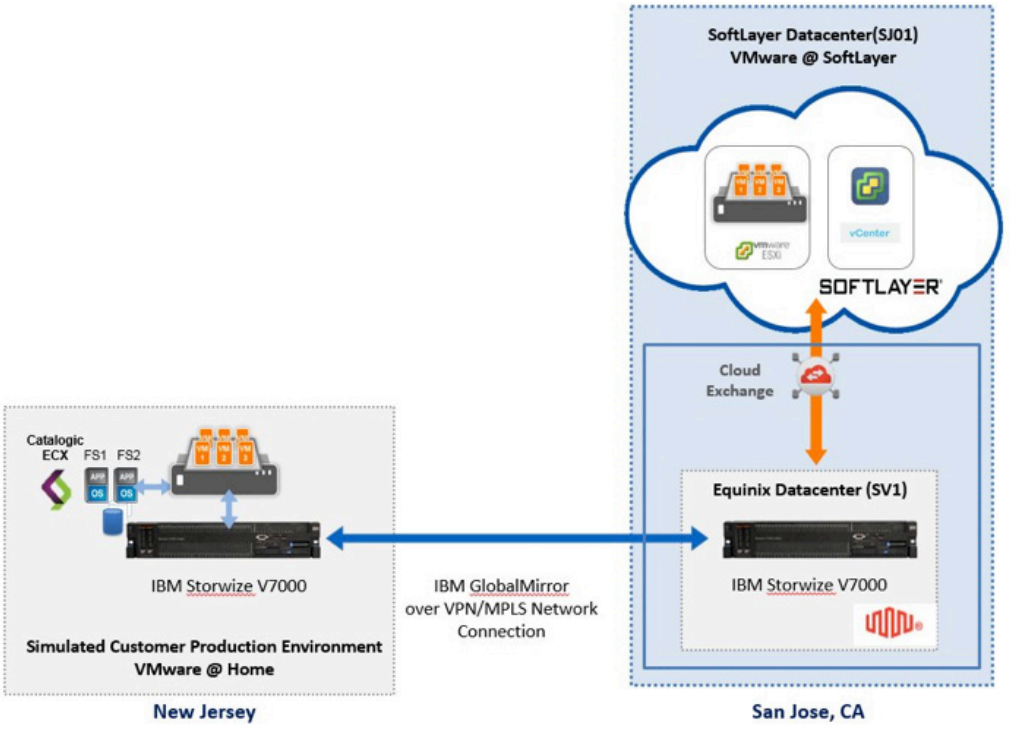

The testing and the results validate the ability to simply and predictably deliver an automated DR process in the SoftLayer IaaS cloud infrastructure. Using the results from this testing and contrasting them to a traditional DR approach, such as one that would test offsite DR readiness twice per year, we have demonstrated how this use case can drive as much as a three times CapEx savings and 200 staff hours in OpEx savings.

Using Copy Data Management and Snapshots for the Hybrid Cloud

Using application consistent array snapshots and mirrors in the hybrid cloud provides a superior disaster recovery solution.

It is useful to understand the details of how a copy data management platform can drive operations.

IBM FlashCopy® snapshots are a space efficient way to capture point-in-time data images that can be leveraged for a number of use cases. Combined with Global Mirror replication capabilities, the Storwize platform lets you quickly and easily move data in and out of different locations, which is essential to being able to extract greater value from a hybrid cloud strategy.

The Catalogic ECX copy data management platform automates and orchestrates FlashCopy and Global Mirror functionality. This powerful and easy-to-use solution replaces the need for scripting and the mix of various tools that today are required to take advantage of the hybrid cloud. With a few clicks, ECX

enables policy-driven workflows to be built on top of the IBM and VMware APIs that leverage the hybrid cloud.

ECX workflows define:

- Which snap copies are used
- Where they are located
- How often they are refreshed
- Who has access to them
- How long they are left in place (retention)

Snapshots can be mounted to servers, promoted to production, or properly torn down to minimize any orphaned or forgotten resources that were being used temporarily. This helps lower operational costs and prevents snapshot duplication, which in turn reduces complexity, management overhead and costs in a hybrid cloud environment.

As an additional benefit, full compliance is achieved since the ECX catalog records:

- All snapshots taken
- All mirror relationships
- Which snapshots are deployed against which applications
- Who has accessed those snapshots

With this solution, the IT organization has a common platform that all system administrators can use to generate the same automated, repeatable and auditable method to leverage copies of data in the hybrid cloud for different business purposes.

Traditional DR Testing

To put a typical disaster recovery workflow that doesn't use ECX into more perspective, here is a granular breakdown of the steps involved.

- Pre-synchronize the storage resources between the production and recovery sites
- Shut down the virtual machines at the production site and prepare them for migration to the recovery site
- Synchronize storage resources between production and recovery sites
- Suspend non-critical virtual machines at the recovery site to make room for migrated virtual machines from the production site
- Change storage access at the recovery site to "writable"
- Power on the virtual machines at the recovery site (these virtual machines can be powered on in a pre-determined priority order if necessary)

For a customer of similar size to our test environment, DR testing is a process that typically takes a full weekend and is administered using many scripts and/or runbook procedures, both of which must be continually updated. Very often recovery tests will fail and significant IT work is needed to diagnose and correct the problem.

IT teams know that testing more frequently leads to greater success, but since DR testing requires such a monumental effort, identifying the best way to tackle "how" to make that happen is a challenge. Moreover, the overall DR operation can be one of the most costly responsibilities for IT and typically requires a second data center. With IT budgets constantly under pressure, painful tradeoffs are often made that negatively influence the overall effectiveness of the DR strategy.

Having the ability to quickly iterate through challenges and fix them as they come up is key to ensuring a successful disaster recovery environment. By testing in SoftLayer and leveraging IBM's storage services, IT has the ability to test, fix, and test again, iterating through the process quickly as they pay only for the time the services are running. Clients can streamline the DR testing process and perform it daily, something that would be unheard of in the traditional style of DR testing.

Modern DR testing can also leverage the different geographic locations of SoftLayer's infrastructure. By making use of different locations, an organization can avoid being affected by large, regional events such as hurricanes. This in turn improves the effectiveness of the company's disaster recovery strategy.

Results Summary

The test results show how Catalogic ECX was able to automate and orchestrate an instantiation of application consistent Mirrors in SoftLayer, taking full advantage of IBM for the automated DR use case. For the DR testing, we used a subset of the simulated client environment. We focused this testing on seven VMs across a range of applications:

- Five VMs with their data stored on the IBM V7000 system:
  - Three application servers: Exchange, Oracle, and SQL
  - Two IT admin servers: Active Directory and a webserver
- Two VMs were IT administrator servers that stored their data on non-IBM storage. For these, ECX orchestrated the creation of VM image copies and replicated them to the IBM platform through the vCenter API as described Section 4.

With the simulated production environment up and running, establishing DR connectivity to the Equinix and SoftLayer environments was quick and easy. The first VM replication jobs began running within two hours of the beginning of our effort to establish the hybrid cloud environment. The more significant time saving came from ECX's automated process of instantiating VMs at the DR site and validating their readiness. When compared to the traditional "brute force" method of testing DR, we determined that the operational savings were quite significant. For a customer that is similar in size to our simulated environment, the traditional process would typically require upwards of 200 staff hours per year. This assumes a DR test effort twice per year, with multiple people spending the weekend planning and executing each DR test.

In contrast, the Catalogic ECX-powered hybrid cloud solution took one person less than an hour to execute. Our tests also validated DR readiness on a daily basis.

In addition to the OpEx savings from the significant reduction in labor to plan and manage the DR testing, the test results also allowed us to conduct a financial analysis that compares the CapEx savings of a "traditional" DR approach to one that is powered by the joint SoftLayer-Catalogic-IBM solution. This analysis demonstrates a 3X cost delta between the automated daily DR approach to that of the status quo.

Financial Savings

The results of our testing allow us to conduct a detailed total cost of ownership (TCO) financial analysis that contrasts the joint ECX-IBM-SoftLayer solution to that of a traditional approach to DR. Our comparison shows cost savings of over $600,000 for a customer of similar size to our test environment, a 3X improvement in TCO over the traditional approach over a 3-year period (shown below in Table 1). While every customer environment and cost structure is different, similar results should be attainable by most customers that would look to deploy automated DR in the hybrid cloud as described in this report, in contrast to a traditional DR operation enabled by a second physical data center.

The TCO analysis below quantifies the various "hard dollar" operating and capital expenditures required to run the operation for both scenarios. While the quantification of the significant staff time savings detailed in sections 7 and 8 would translate to significant operating expenditures (OpEx) savings, it is not included below. By taking advantage of SoftLayer and the flexible compute resources that can be spun up and spun down on demand, the cost savings are dramatic when compared to a second physical DR site. Included in our costs of $2.765/day for the hybrid cloud compute environment is the cost of the building space, the cost of power and cooling, the cost of the hardware for running our tests, as well as the cost of software. Our testing demonstrated the usage of the SoftLayer cloud for a maximum of two hours per day (often much less). Based on our testing, we projected out for three years and conservatively assumed an average two hours per day required to test DR every workday for the modeled environment, the total three year cost would be a little over $4,000 / year. The total breakdown of the savings is in Table 1 below and demonstrates the joint SoftLayer-Catalogic-IBM solution to be more than 3X less expensive.

As the testing and financial analysis demonstrate, by combining the power of SoftLayer, the efficiency of IBM snapshotting capabilities, and the automated orchestration of Catalogic ECX, clients can expedite their journey to the hybrid cloud to test their DR operation, save a great deal of money, and build a more agile business in the process.

Table 1: Total Cost of Ownership Savings Leveraging the Hybrid Cloud

Cost Category	Owned DR site, no ECX		ECX & SoftLayer Cloud		ECX & SoftLayer Cloud Cost Advantage
	Cost	Details / Assumptions	Cost	Details / Assumptions	
Datacenter building	$720,000	$20/sqft /mo x 1000 sqft = $20k/mo x 36 mo	NA	NA, part of SoftLayer cloud costs	$720,000
Power/Cooling	$10,000	$.11/kWh/device x 8760 hrs/year x 10 devices	NA	NA, part of SoftLayer cloud costs	$10,000
Hardware: Storage	$50,000	3 year cost = $2500/TB x 20 TB	$50,000	3 year cost = $2500/TB x 20 TB	$0
Hardware: Servers	$10,000	3 year cost = $5000/server x 2 servers	NA	NA, part of SoftLayer cloud costs	$10,000
Hardware: Networking	$20,000	2 Cisco switches @ $10,000/switch	NA	NA, part of SoftLayer cloud costs	$20,000
Software: VMware	$5,000	VMware vCenter server, standard price	NA	NA, part of SoftLayer cloud costs	$5,000
Software: Storage (snapshot)	$31,000	All software included	$31,000	All software included	$0
Software: Storage (replication)	NA	All software included	NA	All software included	$0
Bandwidth	$135,000	100 mbps bandwidth = $3,750/mo x 36 months	$72,000	100 mbps on 1 gig Level 3 line = $2,000/mo x 36 months	$63,000
SoftLayer Cloud	NA	NA	$4,313	$2.765/hr x 2h/day x 260 days/yr x 3 yr	($4,313)
SoftLayer Cloud Connecdt	NA	NA	14,364	$399/mo x 3 yr	($14,363)
Equinix	NA	NA	$90,000	$3000/mo x 3 yr = $90,000	($90,000)
ECX license + year maintenance	NA		$25,000		($25,000)
TOTALS	$981,000		**$286,677**		$694,323

Conclusion

This validation report shows that by utilizing the combination of SoftLayer, IBM and Catalogic, IT has a new, more powerful way of harnessing the value of the hybrid cloud, without adding complexity or requiring additional expertise. In fact, this solution dramatically simplifies operations and reduces the required labor to manage the flow of data to and from the cloud. The ECX copy data management software platform, in conjunction with IBM SVC/Storwize and SoftLayer, can automate and orchestrate all of the IT organization's key operations that rely on copy data such as automated DR, test/dev or DevOps and business analytics. We showed that the automation of a DR solution provides clients with:

- The ability to take advantage of flexible hybrid cloud resources, lowering total costs by more than 3X
- An automated process for getting the right data to the cloud, saving hundreds of IT hours per year in DR validation testing
- The ability to move data into the cloud quickly–overcoming one of the primary obstacles of cloud adoption today

- The ability to orchestrate data copies in the cloud for multiple use cases beyond just automated DR testing
- Visibility into what data is in the cloud
- The operational control of copy data across the organization, including the cloud, thus ensuring that the IT team is only using (and paying for) the resources that are needed

The impact of Copy Data Management in the hybrid cloud is transformational, and as this validation report shows, implementation is really quite simple. Having a Copy Data Management solution such as ECX is the best way to understand where snapped and mirrored data lives, allowing IT to better utilize those copies for business operations like automated DR. IBM SVC/Storwize allows IT to utilize the same processes that they use today to create a seamless, secure process for leveraging data into the cloud. SoftLayer provides on-demand ECX compute resources, and eliminates the need for the client to own a physical DR environment.

Moreover, other business functions that rely on copy data become simplified and less expensive to manage. A proper copy data management solution, such as Catalogic ECX, enables greater data access for use cases such as test/dev or business analytics, and reduces the burden on IT to keep these environments fresh through automation. The data can be automatically refreshed as often as a snapshot is taken.

As IT looks for solutions to reduce operational and capital costs, while improving its ability to deliver mission critical services to the business, use of the hybrid cloud holds tremendous potential. This paper demonstrates that using IBM SVC/Storwize and Catalogic ECX with SoftLayer overcomes many of the barriers to cloud adoption today, allowing IT to provide simplified access to copy data for multiple business operations, thereby saving significant spending while delivering superior value to business.

About SoftLayer, an IBM Company

SoftLayer, an IBM Company, provides cloud infrastructure as a service for customers of all sizes: from web startups to global enterprises worldwide. The company was acquired by IBM in 2013, and is part of IBM's Global Technology Services business unit. Learn more at www.softlayer.com.

About Catalogic Software

Catalogic Software is the leading software provider of instant Copy Data solutions. Data is more critical to business than ever. Business must leverage their data copies to deliver operational efficiencies, while simultaneously delivering greater business agility. The next great leap in Data Management will be Copy Data Management. Our solutions provide for instant visibility of Copy Data or files in an environment and provide instant data access for recovery, Disaster recovery, test/dev, DevOps and business analytics. Learn more at www.catalogicsoftware.com.

50 Tice Boulevard, Suite 110
Woodcliff Lake, NJ 07677
201.930.8200
U.S.A.
catalogicsoftware.com

United Kingdom: +44 (0) 129 380 4734
France: +33 (0) 1 46 67 46 10
Germany: +49 (0) 2102 3076140
Singapore: +65 (0) 9777 1859
Netherlands: + 31 (0) 20 347 23 88