DCIG

Home   **Offline with DCIG**   **Buyer's Guides**   **Interactive Buyer's Guide**   **Products**   **Contact**   **About**

Spotlight Blogs

symantec.

NEC
Empowered by Innovation

EMC²
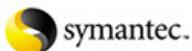
Entry Sponsorship

Ads by Google

Virtual Cloud Servers

Cloud Data Storage

Storage Server

Solution Categories

Archiving (187)
Big Data (32)
Business Continuity (96)
Cloud (41)
Data Management (118)
Data Protection (284)
Data Retention (35)
Deduplication (214)
Disaster Recovery (103)
Electronic Discovery (192)
Information Management (85)
Networked Storage (205)
Physical Tape (55)
Replication (128)
Search (44)
Security (48)
Storage Management (86)
Storage Systems (202)
Tape Systems (34)
Tiered Data Systems (30)
Virtualization (302)

Popular Analysis

Data Center Management 101 Part I (Cable Management)
Email--Not Just Communication But a Legal Document of Record
The SSD Garbage Collection Problem Explained In Depth by WhipTail CTO James Candelaria - Part 1
To RAC or Not to RAC? Veritas Cluster File System Responds "No RAC Needed"

ZL Technologies Brings Virtual Scale-out to Enterprise Archiving and eDiscovery

By **Ben Maas** on March 7, 2012 5:00 AM | **Permalink** | **Leave a comment**

The connection was reset

One of the most exciting and terrifying times in the lifecycle of a company is transitioning from a small to mid-range or mid-range to enterprise sized company. Well led companies that survive those transitions have often been planning for the occasion for some time. The longer they have been planning the more likely they've become aware of the need for long term archiving. *Of everything.*

In some circumstances companies may find that in order meet compliance or litigation standards every instant message, document, revision of a SharePoint archive, blog or email in the company may need to be retained. When viewed from this perspective the volume of data created by even a small company quickly becomes staggering. Worse, if left unclassified all that data becomes next to worthless.

Clearly, the sooner companies in circumstances such as these adopt systems and processes to capture and classify all that data the better off they will be. *The trick is identifying the processes and infrastructure that can scale with these businesses as they grow.*

An entire class of software products has been built to meet this herculean task. One of the most feature rich is **ZL Technologies** flag ship product, **Unified Archive**. Companies with large data retention requirements that are looking to grow should consider adding ZL UA to their list of candidates.

Data stores it can manage range from a small workgroup with a few hundred gigabytes under management to several petabytes. In one **case study** ZL states that a "top 5 US bank" has over 165,000 mailboxes under management.

Deployments of such wildly different scale are made possible by what ZL refers to as "**ZL GRID**." This refers to ZL's utilization of virtualization, clustering and other "cloud computing" techniques. The use of these techniques provides customers with a great deal of flexibility, scalability and performance.

Virtualization promotes flexibility and scalability by allowing the administrator to dynamically add and remove processing resources. Think of the approach as a variation on the old saying "*Many hands make short work.*" The system works by running a number of smaller independent components that break tasks up in to small chunks and then spread the work over several virtual machines.

Because the virtual machines work in a cluster, the work of the cluster as a whole is not interrupted as VMs are added or removed from the cluster. This "cloud" approach allows for permanent long term scaling as a company grows but also enables them to adjust to temporary surges in demand or even one time jobs such as ingesting a large legacy data store.

When demand increases additional virtual machines are brought online to meet demand. As the peak subsides the extra virtual machines can be shutdown, returning their resources for use elsewhere.

A common scenario would be during the "ingestion" of a legacy archival system into the new ZL data store. Ingestion is the process of importing existing data into Unified Archive. This is a very resource intensive process as each document, email or similar structured data is indexed, categorized and deduplicated for storage.

To prevent any user performance impact administrators could temporarily add additional virtual appliances, perhaps on non-production hardware, to handle the load of the ingest process. During a spike in user load the ingestion machines could be suspended and user facing VMs could be powered up. After the user load returns to normal the ingest

Buyers Guides

INTERACTIVE
2012
DCIG
BUYER'S
GUIDE

Search

Search

Newsletter Signup

GO
Enter your email address and select "Go"
Privacy by SafeSubscribe℠

Share and Subscribe

Subscribe to DCIG feed
Follow DCIG on Twitter
ShareThis

Offline with DCIG

DCIG Webinar - Deep Dive into VM Backup Software Features
A Storage Buying Decision Brings Over 20,000 Factors into Play
DCIG and IT Central Station Bridge the "Product Short List - User Opinion" Gulf
DCIG 2012 Backup Appliance Buyer's Guide Press Release
Life is Too Short
DCIG Delivers Analyst Services that Beat the Competition
DCIG Slashes IT Product Selection Risks and Time First Interactive Buyer's Guide for Storage and eDiscovery Buyers
One Hit Wonders
Vital Support Systems Has a Gift Waiting for IT Server and Storage Administrators in the Midwest
A Self-Inflicted Haircut
An Informative Blog is the

DCIG Disclaimer

DCIG writes evaluations of products and services in the storage and electronically stored information (ESI) markets for consumers, public relations firms, business analysts and other interested companies. Our analysis is an informed inside look made possible through business blogging agreements.

VMs could be restored, continuing from where they left off.

Scaling is not just for ingest though. *Each component in ZL's architecture can be scaled separately.* For instance, more VMs can be assigned to the eDiscovery component when legal is creating heavy traffic in the morning. Those VMs can be released when unneeded and then additional VMs spun up for the compliance engine in the afternoon. Clearly, such flexibility to focus resources where they are most needed is a great advantage over legacy systems.

High-availability clustered computing systems such as **ZL GRID** have the further advantage of being very resilient. If one or more nodes are lost the rest of the cluster is able to identify that they are not responding and reprocess the work that was lost. This often occurs without any impact to the user. This failover can even occur over multiple geographic locations.

For instance, in most cases companies will elect to have their most frequently used data stored locally. As the data ages it is moved further away, often to an off-site location or into the semi-public cloud. **ZL GRID** can be configured to failover to a remote site to prevent service interruption, though at reduced performance.

ZL Technologies makes no bones that they are targeting enterprise customers. Unified Archive is one of the most feature rich platforms on the market and not every company can fully utilize its power. However, businesses that have the need and are expecting growth will have a hard time saying no to a system that can scale to meet any challenge thrown at it.

Solution Categories

- **Archiving**
- **Electronic Discovery**
- **eMail Archive**
- **Governance Risk and Compliance**
- **Information Governance**
- **Virtualization**

Rate this entry

Rating: **3**/5 (2 votes cast)

Leave a comment

Optional: Sign in with **Facebook** | **Twitter**
Name

Email Address

URL

☐ Remember personal info?

Comments

Captcha:

Type the characters you see in the picture above.

Preview Submit

Recent Entries

December 2012