Think Global, Act Local!

How to Massively Simplify Data Protection for VMware

Prepared for Cohesity
June 2016

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Backup is still today a major concern to end-users when it comes to IT operations, and the bigger the organization, the greater the issues. The underlying reasons for data protection issues, and for backup issues in particular, still lie in the backup process itself.

Most enterprise backup software was born before the virtualization era. It is complex to operate, with user interfaces and schedulers being all but user friendly. This is a consequence of the complicated software architecture built on agents, servers, media managers, and so forth. Furthermore, the backup process, when not fully integrated with the various software and hardware layers, can be intrusive and impact the performance of the production environment during the back up.

The overall cost of these solutions is high - you need to develop and maintain operational skills, and the software licenses are rigid and expensive. With this high complexity and cost, you'd expect these solutions to provide the desired restoration SLA's. In fact, and especially for VMware environments, it is quite unlikely to achieve instant recovery functionality for VMs or for single files stored in them. This current scenario is becoming even less attractive as most solutions cannot offer effective cloud options for additional capacity, e-vaulting, archiving or disaster recovery. This limits your ability to take advantage of the continuously-shrinking price and flexibility of the cloud.

At the same time, the rest of the infrastructure is evolving drastically towards hyperconvergence. Now, with the exponential growth of storage capacity and the number of VMs under management, infrastructure simplification is a must. Simplicity helps to scale and ease IT operations. As a consequence, IT organizations now desire to hire more general purpose system administrators than in the past, professionals who are inclined to standardize on VMware-aware infrastructure, or at least VMware-integrated to simplify their day-to-day job.

Backup, and data protection in general, must follow the same trend as the rest of the infrastructure, where compute and storage are collapsed together and become much more transparent and easy to use.

Backup, and data protection in general, must follow the same trend as the rest of the infrastructure, where compute and storage are collapsed together and become much more transparent and easy to use. Hyperconverged secondary storage has the same characteristics and benefits as hyperconverged infrastructure if implemented correctly, but instead of providing compute resources to the hypervisor, the compute component delivers advanced backup, file services, and analytics.

Cohesity offers a hyperconverged storage solution which can significantly simplify backup operations and reduce costs, while enabling new data protection paradigms and opening up new opportunities in data recovery, cloud-based disaster recovery and active archiving. By taking advantage of VMware APIs for data protection, and integrating with VMware vCenter, Cohesity can substantially enhance user experience while improving backup operations and restoring speed and granularity. And thanks to the scale-out design, these benefits are available to organizations of all sizes.

HYPERCONVERGED STORAGE AND VMWARE BACKUP

The Cohesity Data Protect platform is not only about data protection, it's much more. It's a comprehensive data platform designed to seamlessly manage all secondary data, exploit its value and make it easily reusable. But, to achieve its goal, backup is one of the preferred ingestion mechanisms and needs to be easy to use as well as transparent and fast when it comes to recovery.

Cohesity Data Protect is a comprehensive data platform designed to seamlessly manage all secondary data, exploit its value and make it easily reusable.

Nowadays, data footprint efficiency is a key aspect for any modern backup platform and storage systems in general. From this point of view, deduplication is taken for granted while other technologies, like space efficient snapshots for example, can make the difference in the granularity of backups and the speed in restores. SnapTree™, a patented snapshot technology available on the Cohesity scale-out File System, is at the core of their solution. Its integration in the backup process allows for as many backups as needed without impacting the performance and capacity resources of the cluster.

Integration plays another important role when it comes to improving efficiency. In fact VMware offers a rich set of APIs that can be leveraged to simplify backup operations. The ability to connect with vCenter, and see all the information about VMs and data stores, allows a complete view of the entire VM environment. This kind of information has enabled companies like Cohesity to rethink the way end users can interact with the system by introducing backup profiles which are now based on SLAs (service level agreements) instead of traditional and cumbersome job scheduling. Consequently, ease of use is vastly improved thanks to better backup automation, while monitoring and alerting are more insightful because they're focused on SLAs and not on single anonymous backup jobs.

Backup is only the first step of the process. Quick data recovery with single object precision is still the most important outcome.

Backup is only the first step of the process. Quick data recovery with precise single object restores is still the most important outcome. Contrary to what happens in traditional backup solutions, hyperconverged storage systems like Cohesity, thanks to clever indexing capabilities, offer a Google-like search functionality which is definitely faster than any other browsing mechanism and scales automatically when adding nodes. When the object has been found, regardless if it's a specific VM or a file within it, fast recovery mechanisms allow direct extraction of the file or spin up the VM directly from a virtual data store exposed by the appliance. It is then possible to vMotion the datastore

back it's original location. In this case, the RTO (Recovery Time Objective) can be measured in seconds. This also allows building superior disaster recovery plans or instantiating as many dev/test environments as needed in minutes, starting from backup images and without impacting production systems.

VMware agentless backup, ease of use and fast recovery features improve day-to-day operations. But, it's not sufficient to achieve the best infrastructure TCO. This is the reason why integration with private and public cloud infrastructure is implemented at the core of the product. Cloud storage, in the form of object stores, can be leveraged to seamlessly manage capacity expansion of the scale-out appliance and take advantage of cloud economics or to implement cloud archiving as well as a cloud-based disaster recovery strategy.

Thanks to the potential of Cohesity analytics workbench, VMware backup will no longer be passive data protection only, but a precious active source of information for the entire organization.

Contrary to backup-only solutions or dumb repositories like VTLs, hyperconverged storage is an integrated system. As such, an embedded and powerful analytics engine makes it possible to crawl data and find useful information or build custom applications to create insightful reports - covering business aspects well beyond the mere infrastructure like, for example, e-discovery, policy compliance and security auditing among others. Cohesity scale-out appliance is intheperfect position to collect and consolidate different types of backups, file services and many other storage needs. Thanks to the potential of Cohesity analytics workbench, VMware backup will no longer be passive data protection only, but a precious active source of information for the entire organization.

In contrast to traditional backup software, Cohesity offers an all-inclusive licensing model which contributes to improve TCO and infrastructure simplification due to reliable cost predictabiliy of the system over time, and with all the functionalities activated. In fact, each single appliance bundle includes all the features offered and is not limited by capacity, number of VMs under protection or application agents. Thus , the end user is free to choose the best data protection strategy without limitations.

BOTTOM LINE

Hyperconverged secondary storage offers the same benefits as hyper converged infrastructure (HCI), but focuses on data management instead of computing. This is why Cohesity hyperconverged solution is a unique platform and addresses a large spectrum of secondary storage needs including backup. Its modern design has enabled the creation of a comprehensive set of innovative features and VMware backup is just one of these.

By collapsing many components of the stack in a single scale-out appliance you can do much more while simplifying processes and operations.

Data protection is a key aspect of any infrastructure. Performing VMware backups correctly, especially in larger environments, can be very challenging, and this is where Cohesity outweighs traditional backup solutions. The major benefits of Cohesity's solution comes from its ease of use, overall efficiency, fast recovery, scalability and cloud integration:

- It is extremely easy to use and powerful at the same time thanks to a policy based scheduler which is focused on SLAs instead of backup jobs;
- Thanks to SnapTree technology it adds efficiency and flexibility to data/copy management;
- Instant recovery options, like file level search/restore in VMs or the ability to spin up VMs directly from backups, reduces RTOs from hours to seconds;
- Scale-out design and appliance form factor do not pose limits to scalability and consolidation;
- Cloud integration means freedom of choice and unmatched flexibility of the cloud while keeping data and costs under control.

By collapsing many components of the stack into a single scale-out appliance you can do much more while simplifying processes and operations. Major operational improvements can be achieved in VMware data protection by following this approach. Futhermore, with Cohesity, your backup infrastructure becomes an asset by extracting even more value from your protected data.

JUKU

WHY JUKU

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help decision makers in their IT environments, to inform and to discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in the IT environment need to be aware of their surroundings: things are changing rapidly and there is a need to be constantly updated, to learn to adapt quickly and to support important decisions - but how? Through our support, our ideas, the result of our daily global interaction on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there - the comparison and the search is global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential in providing an honest and productive vision. That's why we have chosen: "think global, act local" as a payoff for Juku.

THE AUTHOR

Enrico Signoretti is an analyst, trusted advisor and passionate blogger (not necessarily in that order). He has been immersed in IT environments for over 20 years. His career began with Assembler in the second half of the 80′s before moving on to UNIX platforms until now when he joined the "Cloudland". During these years his job has changed from highly technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm deeply focused on supporting end users, vendors and third parties in the development of their IT infrastructure strategies. He keeps a vigil eye on how the market evolves and is constantly on the lookout for newideas and innovative solutions. You can find Enrico's social profiles here: http://about.me/esignoretti