

White Paper

ZFS

Introduction

The IT industry is developing by leaps and bounds, being driven by unsatiated consumer demands and trends. This causes an influx of data at a never-before-seen speed and amount. Along with the incredible growth of the amount of data, its value grows just as fast. The usual way of dealing with such issues would be to pile up costly hardware storage infrastructure components. Commodity hardware wouldn't cut it, so enterprise-grade proprietary storage arrays would have to be bought. This would be the usual way of accounting for necessary capacity, growth prospects, and data safety. This, in turn, leads to unobjectively and unnecessarily high implementation expenditure, as well as the increased total cost of ownership (TCO), which is not something Enterprise ROBO, SMB or Edge deployments, are looking for.

Problem

Contemporary businesses of all sizes must store enormous amounts of data while ensuring its safety. The modern demand for storage systems is that they are to be operable round-the-clock and maintain constant data availability — no system failures are acceptable regardless of whether one node fails or more. Apart from this, storage systems have to be resistant to the failure of local drives to uphold constant application uptime, which means that the whole IT environment has to maintain data integrity at all costs. With this in mind, companies have to build their storage infrastructure with a file system in mind that is able to store huge amounts of data while keeping it safe from corruption, highly available 24/7/365, and at a reasonable cost.

Solution

StarWind coupled with ZFS provides customers with an Enterprise level storage system that enables unmatched performance, solid data integrity, and multi-level local redundancy at SMB cost. Protecting an IT environment from any failures, including node failures, is an innate feature of StarWind. StarWind's ability to symbiotically integrate with ZFS enables local redundancy with any RAIDZ level of choice ("mirroring," Z1, Z2, and Z3). On another note, ZFS makes data and metadata integrity checks, further enhancing the protection from silent data corruption. It is also possible to schedule disk scrubbing to avoid potential bit rot errors. Integration with ZFS also ensures the ease of scalability for prospective growth as the file system supports up to 256 trillion yobibytes (2^{128} bytes) of storage capacity! Thus, all data on every StarWind-based cluster's layer is protected, and each node/disk/file system is tirelessly monitored for health.

Conclusion

The combination of StarWind and ZFS within a VSA (Virtual Storage Appliance) enables organizations of any size to acquire an Enterprise-grade storage system at a reasonable SMB price. The smooth interoperability of these two resoundingly increases data protection levels, accounts for constant maximum application uptime, and virtually eliminates any chances of data corruption. The integration of StarWind with ZFS is a definition of a system that is always prepared for any disasters, functions on its best behavior, and prepared for any potential business and infrastructure growth demands.

Contacts

**Phone number**

+1 617 449 77 17
+44 2037 691 857
+34 629 03 07 17

**Fax**

+1 617 507 58 45

**Voice Mail**

+1 866 790 26 46

**Customer Support Portal**

www.starwind.com/support

Support Forum

www.starwind.com/forums

Sales

sales@starwind.com

General Information

info@starwind.com

To learn more, visit www.starwind.com