

SCI StorInt™ Dispatch - EMC070904 Announcing: Avamar Data Store and Virtual Edition

This Silverton Consulting (SCI) Storage Intelligence (StorInt™) Dispatch provides a summary of EMC's recent Avamar system announcements providing a new rack packaged server product and a new VMware specific application package for Avamar technology.

EMC Avamar technology

Avamar technology is a complete end to end backup and recovery solution featuring source data de-duplication technology. It can be a complete standalone solution that eliminates the need for traditional backup infrastructure, or it can be used for specific problem backup areas. The technology supports sub-file, intelligent data de-duplication, independent of file type or data base. Avamar technology delivers quick restores from disk. In addition, Avamar technology supports disaster recovery by replicating backup data across the WAN to another Avamar server.

Avamar technology is deployed via client agents and servers. De-duplication is done first and primarily at the client and again at the server. Client agents consult their local and server cache for maximum de-duplication. Server de-duplication is done to provide cross client data de-duplication.

Avamar Data Store announcement summary

Avamar Data Store takes the Avamar technology and packages it into a pre-racked solution containing backup, recovery and replication software and all needed server, switch, and storage hardware. The Data Store comes in a multi-node version and a single node version. The multi-node Data Store comes configured with a spare server and a utility server that supplies reporting and administration capabilities for the rack.

Each node holds 1TB of Avamar de-duplicated backup data which, under a normal system backup schedule with weekly fulls and daily incrementals, would require approximately 35 TB of traditional tape or disk storage. A Data Store can have up to 16 data nodes plus 1 spare and 1 utility node to support almost a half a petabyte of traditional, cumulative customer backups.

Also, multi-node Data Store implements a RAIN (redundant array of inexpensive nodes) architecture, which can sustain a node outage and still deliver backup data without interruption of service or loss of data. RAIN architecture is in addition to normal RAID data storage that each node uses internally to hold and protect customer backup data.

Stored in each RAIN node along with customer backup data is an index to the backup data, which identifies each component of backup data. Because the index is stored with the customer backup data, it allows the Avamar system to scale to very large backup loads.

Avamar Virtual Edition announcement summary

EMC also announced a new Avamar software-only package specifically targeted at VMware environments. The Avamar Virtual Edition is a tightly integrated software solution that operates the Avamar server in a VMware virtual machine. Just like any Avamar server, it provides a backup/recovery target for a number of Avamar clients. This virtualized deployment allows customers to run a full backup, recovery and replication solution on top of VMware Virtual Infrastructure and augments Avamar's ability to backup VMware.

Along with normal Avamar virtual machine clients, Avamar agents can back up at the VMware service console level or can do backups at the VMware consolidated backup server. This provides multiple client configurations to support VMware complexes. Avamar file data client agents are free, but Avamar database agents are separately charged items. Pricing of database agents or servers was not disclosed.

In VMware environments normal backup processes are processor, disk and network I/O intensive. With Avamar, one can reduce backup network I/O significantly (by 90%), and even though data de-duplication is done mostly at the client, backup processor utilization can also be reduced significantly. EMC stated that processor utilization for backup activity can be reduced by 80-85% once Avamar is in place and active.

Announcement significance

The pre-configured Avamar Data Store offers an option for Avamar technology that is easy to deploy and configure. With the exception of client agents, everything you need for backup, recovery and replication is packaged as a single solution and can be readily scaled to meet very large backup needs. In addition, the fact that it is easy to install and deploy makes selling the product through the channel more productive.

Avamar Virtual Edition provides increasing integration and usability in a VMware environment. As VMware has taken over server virtualization, anything which eases deployment and use in this environment addresses a tidal wave of business.

De-duplication can be done at many points along the data protection activity chain. The main advantage of doing it at the source is that it reduces all the data processing and throughput needs from that point on, obtaining the technology's advantage across the whole chain of data protection processing. In green field environments, deploying Avamar from the ground up makes sense. In larger environments with existing backup infrastructure, Avamar should be deployed for remote offices, VMware and large NAS environments, where source de-duplication provides many benefits. The packaged solution might appeal more in the SMB space, since everything is packaged in one server.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in Colorado, USA offering products and services to the data storage community