Innovation

Discover the technology of the DDP
Why DDPs way differs from others

The shared storage solution of the DDP - Dynamic Drive Pool - the Ethernet SAN system differs from all others. We want to invite you to join a conversation with Jan de Wit, CEO of Ardis Technologies, telling you more about the DDP technology and its benefits.

Jan, we often spoke about what's unique with DDP systems. DDP is an Ethernet SAN system using *i*SCSI. Could one say that the basis for all the unique features is the Ardis Virtual File System AVFS, the metadata controller developed by Ardis Technologies?

Jan: The Ardis Virtual File System AVFS is the technology behind the DDP, yes. It manages access to the Ethernet SAN and enables that one or even more DDPs - if you have a Cluster configuration - are in One Namespace. That's why there is only one file system, the "ddpvolume", that holds all files, folders and FolderVolumes.

The DDP FolderVolumes are surely one of the main features that are unique to the DDP. What's making them so special?

Jan: With the DDP technology, the metadata (the virtual file system) is independent of the data. The data itself is stored in Data Locations. The metadata show up as one virtual volume ("ddpvolume"). This "ddpvolume" can hold folders with volume properties – the so-called FolderVolumes. For each FolderVolume you have a precise access control, you can assign quota to manage capacity and they also enable a unique way of using SSD Caching and deciding dynamically where the data resides by selecting Data Locations.

> *The DDP has only one file system, the "ddpvolume",*
>
> *that holds all files, folders and FolderVolumes.*

You just mentioned one of the nowadays "buzz words": SSD Caching. I remember that you told me that what DDP means with caching is not very much related to the caching methods used by other manufacturers.

Jan: That's true. Most people confuse metadata SSD caching with data SSD caching. Very often SSD Caching is used as a synonym for caching of metadata. For us, the main topic is the data SSD caching. It's another big advantage of the DDP technology that although iSCSI itself is a block-based technology, the DDP systems work file based thanks to the AVFS file system. This enables file-based caching as well as advanced caching modes. You have full control of the SSD Cache and you can configure it for each of your FolderVolumes to increase performance where and when it is needed.

Jan de Wit, CEO Ardis Technologies

Discover the technology of the DDP
Why DDPs way differs from others

*DDP stack with 3 Data Locations
and Connect Screen*

Okay, that sounds a bit abstract to me now. Could you give me an example?

Jan: Let's say you have a DDP system with 3 Data Locations: "Cache", "HD1" and "HD2". "Cache" consists of SSD drives, HD1 and HD2 consist of normal spinning hard drives. To start your project you create a FolderVolume to hold your project data. It is a virtual volume so you can set the caching method and the Data Location specifically for this FolderVolume. Choose "On-Demand" as cache method and "balanced" as Data Location. All data that you now copy/ingest into your FolderVolume will be distributed Load Balanced to HD1 and HD2. This Load Balancing already increases performance. Plus all data in use will be also copied to the SSD Cache, which again increases performance. For projects with very high-performance requirements you also can use the SSDs as primary storage, so all files are directly stored on the SSD cache.
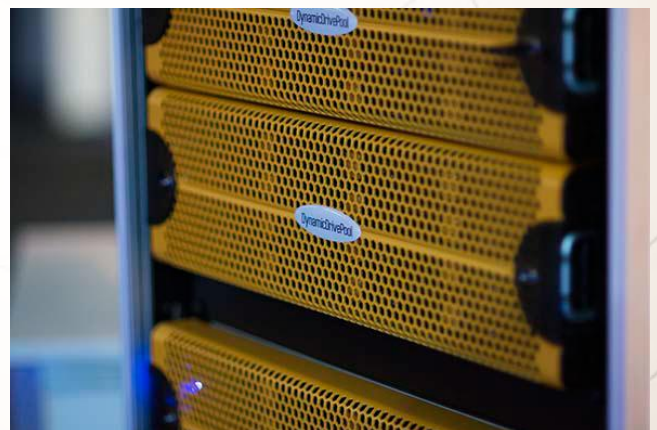
The example "DDP system with 3 Data Locations" sounds like one single DDP system. How does the procedure differ when the setup is a Cluster configuration?

Jan: There is no difference as it is all in One Namespace. As there is always just one "ddpvolume", the virtual file system can hold 1 DDP or multiple DDPs. The 3 Data Locations mentioned above can be all in one single DDP system, but they could be

The DDPs can be of different types and build dates and even do not have to be in the same room or city. That's why DDP gives a real Scale Out solution: Linear Scaling in bandwidth and capacity, desktops have Parallel Data Access to the DDPs in a Cluster, and you can easily add more DDPs without changes to the directory. also in 3 different DDP systems/cluster nodes.

That sounds to me that especially with a DDP Cluster configuration you are pretty safe for the future?

Jan: Well, if you realize that you can add new DDPs or remove old DDPs any time without interrupting the workflow, if you realize that you easily can increase capacity and performance if necessary, if you realize that nevertheless, administration is always for one virtual volume only: yes, that means that a DDP Cluster is able to live forever and to stay always up to date.

DDP stack detail