

## Overview 1

Now, this first module is not my favorite because it contains a lot of slides and a well, I can't get around it because I have to do an overview first. I hope you bear with me in this one. So let's start. Let's first have a look at some possible configurations solutions that Netapp offers. There's a single node cluster, which is one node for me, a cluster all by itself. There's an ha pair, which means that you have two nodes that are combined and interconnected and share the same storage in which can be scaled out by adding more ha pairs to it to generate more storage and better performance. There's a FlexRay which doesn't have physical storage itself, but gets London's from third party storage. There's ontop select, which is a software defined storage solution and Metrocluster, which is a pretty expensive solution that needs an additional 100% of physical storage in order to be able to mirror the data from one data center to another data center in sync.

Now before we look into these configurations, we have to look at the note itself. So I'm not talking about the single node cluster here and talking about the physical node that is part of a single node cluster or is part of a multi multiunit cluster like any other computer. Uh, the NETAPP solution has something which is called Rem, which is your memory and CPU, which is pretty obvious. What's different with a lot of other solutions is that there is something which is called compact flash and a compact flash, a store the operating system. In many environments, the operating system is on some disc or on some petition, but with Netapp is in the compact flash, which is on the motherboard and can contain up to two images of the same operating system with different versions. Obviously this happens when you do an upgrade. Then your comic flesh will contain two images of the operating system.

At boot time, the loader will look in the compact flash and find the kernel loaded and the rest of the boot procedure we'll continue. The next thing which is also important is the service processor. The service processor will offer you access to the system itself, even if your system is down so it has its own IP address and you can connect to the service processor and for example, create core dumps or power off or power on the system. Well, whatever. So there's a number of remote capabilities you can execute when you access the service. Persistent and other vital part of the NETAPP environment is on the node. You have something which is called MV Ram or envy, ma'am. The and envy men conceptually are the same. And throughout this training, whenever I have to talk about this, I will refer to it as envy rem, but I also mean MP men.

When I say I'd be rent, I'm not trying to be funny, but it's, it's, it's more or less the same. What have you, ram is used for is the following thing. When a client writes data that the data will enter in the rem environment. Okay. Now you cannot as a, as a node, you cannot say that the data was written to disk because the data is not on. This kit is it's only in rem and if you would lose your note or the note would be reset, then the client would lose its data. Now Netapp is thought of a very pretty solution. What they do is the minute the block of data enters rem, it will be copied to envy ram instantaneously and the minute it is in MV Ram which is battery backed, the client can be acknowledged that the data was written so the client can continue and uh can rest assured that the data is safe and eventually we'll be on disk.

This is all taken care of by something which is called the creation of a consistency point. So based on particular triggers, the data in memory will be written to disk after some time or because MV ram is filled up or for multiple other reasons. The minute the data is on disk is physically written to the disk drives then and we ran will be flushed and that consistency point is complete. What is meant by consistency point is that you have a new consistent view of the file system on disk. We'll go into that later on in another module. For now, you have to know that Andrew Ram is there just to speed up the acknowledgement to the clients and still make sure that your data is safe.

Now let's start with a single node cluster. What you see in this picture is a single node and three dish shelves. This means that if you lose the node, you lose access to the dish shelves. Um, also what you see is a data network and a management network. Those are the only two networks you need in a single node cluster. Now first of all, why is this called a cluster? This is called a cluster because of the fact that you run the cluster software. And also if you want, you can add a second node to the same cluster on the fly. Doesn't mean that your clients experience any downtime whatsoever. So you can on the fly add a node to the cluster and remove the note if you want to. Now, why would you run a single node cluster? Because of the fact that you have no availability?

Um, if the node dies, your data is no longer accessible, well maybe you want to run this solution as a backup destination. So you have a multi node cluster and you want to back up your production data. You can do that to a, um, to a cluster that is not necessarily up all the time. That would be one reason why you want to run a single note. And of course it doesn't as much as when you buy two nodes or even four nodes. So the advantages pricing and the disadvantages availability, uh, the networking. Um, just one more remark about the networking at the best practices that you run. The data network and management network network across different, uh, physical switches. Um, but if you want to, Netapp does support that you run both networks across the same physical switches, but again, it's not the best practice.

Now if you add a second note to the single node cluster, you automatically create an ha pair. What we'll be striking about this is that all of a sudden you see that you've not only got two data network and a management network, but you've also got something which is called the cluster interconnect. And you have something which is called the HPA interconnect. Now, the Ahj to connect is used for an dram mirroring. Remember when we were talking about one note about the configuration of a node, that there is something which is called MV Ram. And the NP ram is used to, um, to enable the controller to say that the data was written even though it's not on discount. Now, what happens if you have an ha pair, it will mirror the NPRM contents from one node to the other node and vice versa. So usually this is done via the back plane, but that's only the case if you have two physical controllers in one chassis, uh, this is something you should think of when you go and do the exam because that may be a question because you don't have to cable the Han to connect unless your controllers are in two different Chessie's.

So again, two controllers in one single chassis do not need cabling. If you placed the controllers into different Chessie's, you do need a separate cabling for the Aha to connect and it's primarily used for and if he ran mirroring. Um, and other important network is the cluster interconnect. This is used for an awful lot of purposes. It is used for a secret using the configuration in the cluster across all of the nodes. It's used for sending data from the volumes to from one day to another because a client asks a data from a node which does not host the volume, then the data will have to be a retreat from

Donna that does host the volume. It's also used for heartbeats and is used for snap mirroring in track clusters. So if your snap snapmirror relationship is inside the same cluster, the data between the two volumes will be sent across the cluster interconnect.

Uh, this needs a minimum of two 10 gig interfaces per node, uh, to build. So you can add, you can add other interfaces, other ports to the same cluster interconnect environment. But the minimum is two ports per node to set up the cluster interconnect. Another thing, of course, that's important if you have an ha pair, is that you share the storage. In this example, you see three does Chels and all shales are connected to both nodes. So if one node fails, the other node is able to take over the aggregates that were originally hosted by the node that has now failed. So let's take a small break and then continue with the overview.