

THE PROTECTION ARCHITECTURE

BY RANDY CHALFANT

Among the priorities for efficient storage management is an appropriate protection architecture.

We all know how important it is to protect the crown jewels of any business, the data. The question is, do we protect it in ways that make sense. Because the typical guardians of storage administration are typically too few and certainly overstretched with things to do day to day, an appropriate data protection architecture is not given the time or consideration it needs and deserves. I know, that probably brings up howls of injustice from those that implement protection schemes, but there is more to the story of an appropriate protection architecture than just throwing capacity at it in the form of backup copies.

Key to understanding an appropriate protection architecture is understand the value of what you are trying to protect. Agreed by analysts and great minds from clinical offices far from data centers – not everything deserves the highest level of protection money can buy. That seems reasonable and most of us would agree. So why is it then, that in countless data centers I have seen the inverse? Nearly everything in them is the best money can buy. Once again, it seems that people just don't have the time to think these things through, so many adopt a one size fits all. There are also unscrupulous vendors that are happy to oversell capacity.

If you wanted to get it right, where would you start?

Absolutely without a doubt, you have to have a scheme to characterize the value of what you are trying to protect, against



the cost required to protect it. There are two major categories to consider. You must first understand the economic value that a contributing application offers to a business when it is running, and then you must understand the impact to a business if that application stops. They aren't the same. The economic impact is typically much larger over time when unavailable. When you know these numbers, it provides an objective basis to rationalize the cost to sustain the value of an application when running, as well as the cost to protect it when it is not. Different applications will have different values, but trends will begin to emerge, and that is the time to assign applications a class of service that justifies the protection cost, and back that with a reference architecture that contains cost to a known technological approach.

Figure 1 provides a basic classification scheme that can be used to judge the priority assigned to data protection and recovery.

This should be the first step of building a protection Architecture. Once you have all of your applications classified, it is an appropriate time to begin to think about how you are going to apply technology based solutions to meet the needs of protection.

Figure 2 provides a map used to see the big picture of protection. There are many things shown on this map, so let's explore it. The orientation of the horizontal or X axis of the chart can be looked at as the performance requirement for recovery. It starts in the one week range, and extends out to less than a minute of time required to restart a failed mission critical application. The vertical or Y axis represents the quantity of data typically found in a data center. There are cones that rise from the recovery time objectives to describe some of the more common protection solutions used. As an example, a level one-protection scheme would represent the smallest amount of data in a data

Figure 1: CLASSIFICATION DRIVES INFRASTRUCTURE

Protection Classifications	Protection Level	Availability Objectives	RPO	RTO
<ul style="list-style-type: none"> Mission Critical Data Most valuable to an enterprise, high access High performance, high availability, near zero downtime, highest cost 	1	99.999%	1 Minute	1.5 Minutes
<ul style="list-style-type: none"> Business Critical Data Important to the enterprise, average cost Reasonable performance good availability, less than eight-hour recovery 	2	99.999%	10 Minutes	15 Minutes
<ul style="list-style-type: none"> Accessible Online Data Cost sensitive, low access, large volumes Online performance, high availability, less than eight hours of recovery 	3	99.99%	2 Hours	2 Hours
<ul style="list-style-type: none"> Nearline Data Cost sensitive, low access, large volumes Less than one-hour access time, automated retrieval 	4	99.9%	1 Day	1 Day
<ul style="list-style-type: none"> Offline Data Archived data, backup or compliance related Very cost sensitive, limited access, ~72-hour seek time 	5	Offline	1 Week	1 Week

center that cost justifies a solution using either a duplicated site, or a RAID 1 or mirrored copy for recovery. Whereas a level four-protection scheme could use anything from corruption copies on disk to backup copies on tape, or virtual tape. What is important about this view is the idea that you must first map the value of data, which gives you a recover point and time objective. That can be mapped into a solution capability, that can then be used to choose the right technology to provide an effective protection architecture, and an effective business rationalization plan.

What is a far more common approach is quite ad hoc by design. Here's how the protection scenario more commonly goes.

Somebody decides they want to create a new application. The application is coded and tested. With that complete, it is realized that a protection strategy must be put in place. Perhaps the initial view for how much storage required for the application is 1TB. This is the point that the call usually goes out to the disk vendor.

Consider the typical disk only vendor's sales approach to this. The vendor approaches with a story of how they are going to save you money. That's always interesting, right? Everybody seems to have the same goal, increase the infrastructure to keep up with the demands of the business, with less budget and

fewer people. No wonder a customer will react favorably to the idea of saving money. The vendor tells you they have the ability to consolidate the many primary storage systems that you currently have down to a few. They explain the super performance and scalability capabilities they have. Because of that, they can take the various data bases and applications that are spread among multiple subsystems and reduce that to one. You get improvements in performance, fewer components to manage etc. You say, OK

that sounds interesting, but why would I put all of my eggs into one basket? They tell you about all the redundant parts etc, and then they say that even with that, if there is a failure, they have a RAID 1 mirrored physical copy to protect you from any failure of the hardware. So far a very exciting prospect, save money by consolidating, protected with a RAID 1. Figure 3 illustrates three tiers of protection. Tier one is Raid 1 mirroring, and is used for fast restarts for protection level 1 listed in Figure 1.

CONTINUED ON PAGE 14

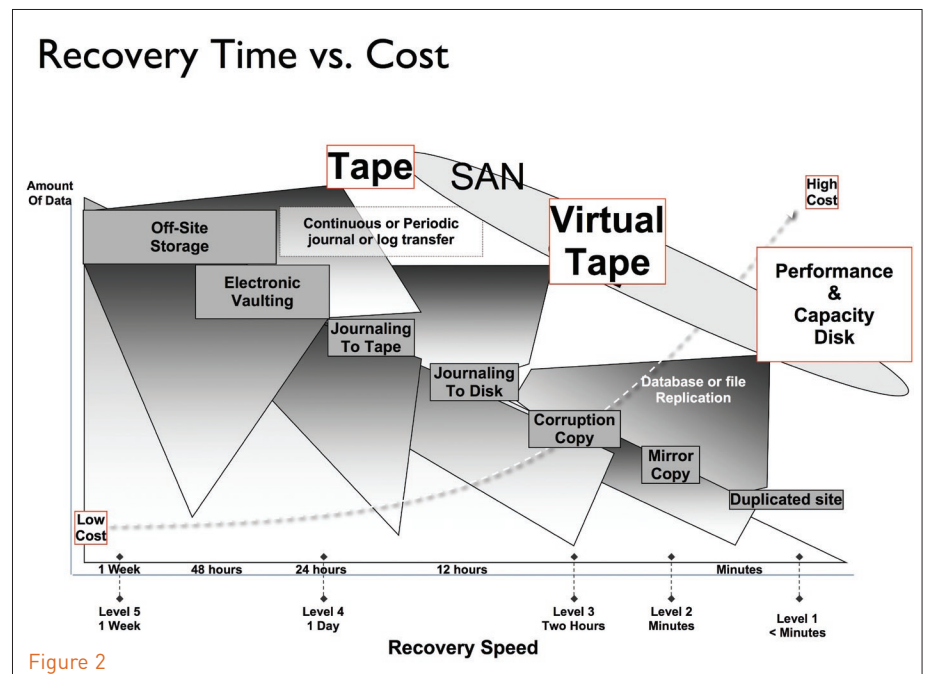


Figure 2

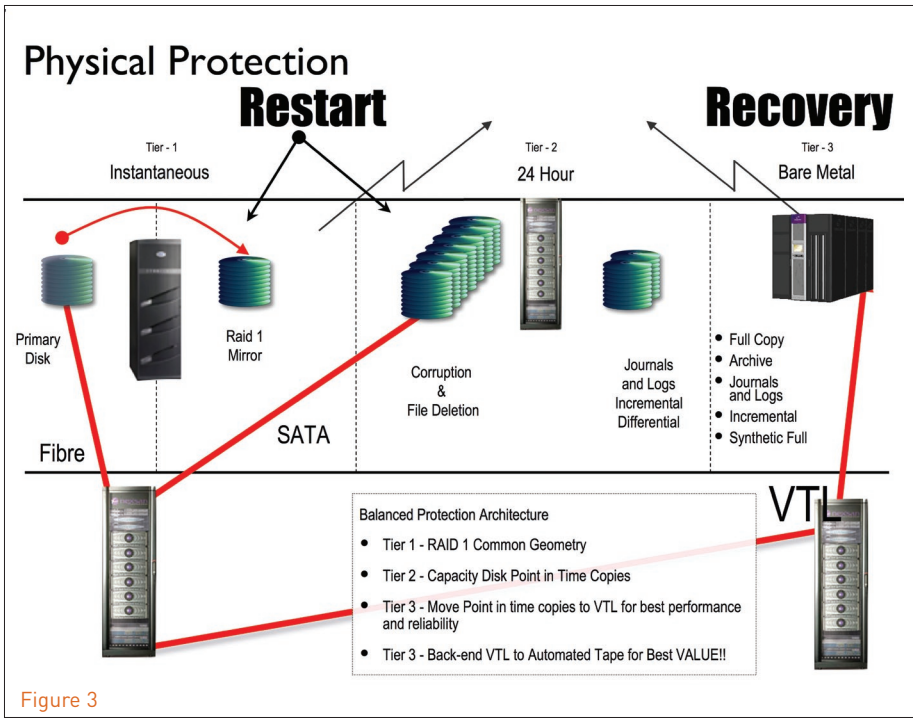


Figure 3

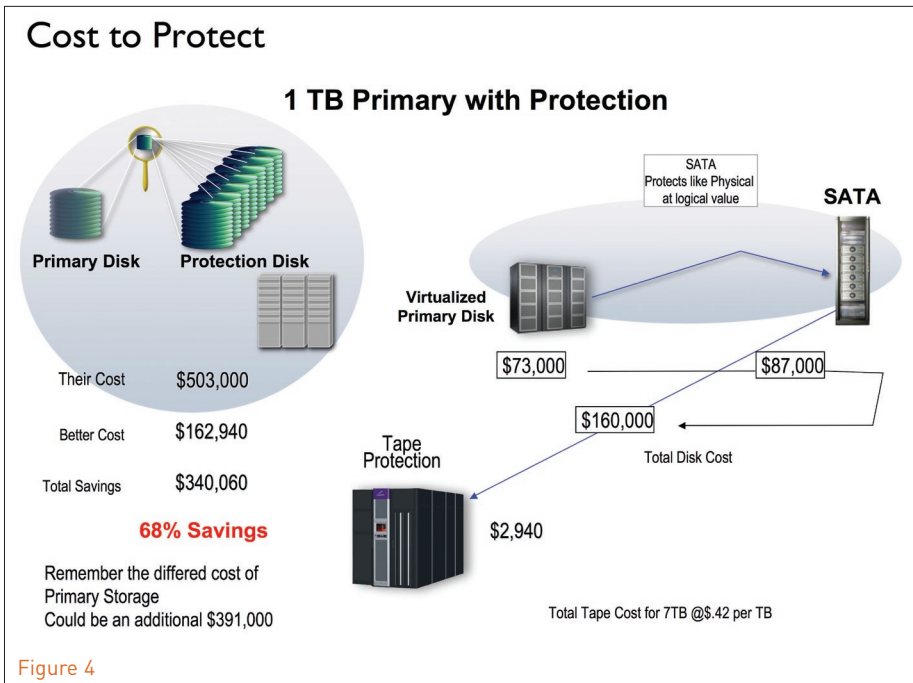


Figure 4

CONTINUED FROM PAGE 13

OK, but what about data corruption? For that, vendors will sell you what can generically be described as point in time copies. Using this strategy, they would have you buy another seven volumes, one that gets broken off every three hours. The value for you is that there are 24 hours of point in time copies to protect you from corruption or file deletions.

Every three hours, you would rotate the current copy out, and re-establish the oldest copy from 24 hours ago.

Hence, if you have a data corruption problem, you simply figure out how long ago the corruption started and recover to a period of time just before that. Right. The problem with corruption is that nearly nobody ever discovers they have a problem within the first 24 hours. One very large web based auc-

tioner learned that the hard way, and guess what, their corruption exceeded the length of time they had mirrored copies of data waiting on redundant disk. They were grateful for their foresight in having copies on tape that did not get corrupted. An important government agency confirmed this problem in a situation they experienced as well.

Even so, if you have class 5, or 5 – 9's of availability application, you will no doubt want to have point in time copies.

I agree that if you need instantaneous restart, then a RAID 1 mirror is the right choice. I agree that corruption protection, file deletion, and journals and logs should be kept on disk for rapid restart for all applications that demand 99.999% (Class 5) or better of availability.

However I know that instead of putting all of this protection on the most expensive tier 1 disk that a tier 2 capacity disk works just as well at a substantially reduced price with the same guarantees for recovery time and recovery point objectives. I also know that Tier 3 tape is and always was designed for a recovery measured in hours or days of time. Somehow, the disk only vendors are trying to reposition in the minds of the market that tape should not be used for a recovery. Look, if you absolutely positively have to have an application back up in minutes, fine – use a high capacity disk, job done. However, don't forget that in the world we live in a plane can run into a building or a dirty bomb can go off. All the redundant copies on disk in the world that are equally destroyed in the same location have no value. If things get really bad and you have to rebuild, vs. restart, forget about disk. You can't afford it on the flat or declining budget you have with an average growth of 30% you have to keep up with. Especially when you consider most customers are wasting up to 70% of everything they buy.

The appropriate approach is to classify the recovery requirements of all data by application value, and then use the right technologies to provide the right protection at the right cost. A balanced mix of infrastructure will keep the point in time copies on capacity disk (not performance disk), and then move

them to a disk masquerading as a tape library, otherwise known as virtual tape, for performance and reliability purposes. By the way, the reliability of the tape is not the issue so long as you are using enterprise class tape, it is the reliability of the environment and that is why you want to use VTL. It helps to mask that. After the captured backup data is safely tucked away on a virtual tape library, the next step is to migrate that to a real tape library for best economy. Contrary to the hype and sensation of the exuberant marketing organizations of disk only companies, tape is still the safest and least costly form of retaining long-term data – period, and it is portable.

The Economics of a Balanced Protection Architecture

Just to give you an idea of how much you can cut out of the cost of a well-balanced protection architecture, let's look at some real numbers and how it maps to a hypothetical example.

Figure 4 shows an environment where we would start with 1TB of protection level 1 data that has a 5 9's of availability reliability objective. If we look at the costs of protection using only tier 1 disk, the costs could be over \$500K. However, looking to the right in Figure 4, we implement a Tier 1 disk, a Tier 2 Capacity Disk, and Tier three tape. This reduces the total costs to \$162,940, saving \$340,060. A 68% reduction in CAPEX alone. This is huge. By the way, using an AutoMaid technology from Nexsan in your tier two disk can save upwards of 70% of the power costs further reducing OPEX. In an economy where people are looking to save every dime they can, while getting more out of the infrastructure already in place, you really need to do things differently than they have been done in the past.

Customer Example

We met a major European retailer with a sophisticated storage infrastructure. Despite this sophistication the organization was again running out of capacity and was close to deciding to install many more terabytes of new primary storage. I met them and we advised them to wait, as I felt they had too much storage already. Sure enough,

erms
Notification and Emergency Management Solutions

- Mass Notification & Targeted Alerts
- Crisis Event Management
- Stakeholder Roll Call
- Online Crisis Document Access
- Hotline
- GIS Mapping
- Personal Emergency Notification

Learn More: www.ermcorp.com • sales@ermcorp.com • 905-829-8216

after completing an assessment audit a few days later we had identified ways to reallocate much of the existing storage capacity in a much more efficient manner, most of which was rebalancing their protection architecture. The end result was to add some virtualization hardware and a realignment of the existing resources to do a better, more efficient job, with an overall savings of about \$3.2M over three years. The incumbent storage vendor wanted to grow them from 32TB to 57TB in three years. We showed them how they could keep up with the current growth rates and shrink at the same time from 32TB to 19TB in three years.

Protecting the value of data is critical, in the world and economy we find ourselves struggling with today, so is protecting your company's ability to meet financial goals. Keeping your infrastructure protected and economically efficient is not only important for the company; it could also mean the difference between having a job, and your company having to let people go.

Conclusions

If you look at this hard, the ability to use a well-balanced protection architecture is based on a strategy to gain value in your business by more efficiently using storage to sustain and to protect your business applications. Yes you must do things a little differently than you have in the past. But in the words of Albert Einstein, "The significant problems we face cannot be solved at the same level of thinking we were at when we created them".

So, may the narrowly focused efforts of business as usual rest in peace, and hopefully you will consider a methodical approach that offers great efficiency as it's reward when protecting your company's treasures.

ABOUT THE AUTHOR

Randy Chalfant has 35 years of experience in engineering and marketing for storage, servers, operating systems, applications, and networking solutions globally. He can be reached at randy.chalfant@mac.com.