

## UNITRENDS DISASTER RECOVERY: VAULTING PERFORMANCE

### EXECUTIVE SUMMARY

The three primary factors governing the implementation of disaster recovery are the characteristics of the network used to transmit the data to the remote location, the amount of data that is specified to be protected and how much that data periodically changes and the software techniques used to effectively manage the bandwidth that is marked for disaster recovery. Each is reviewed in this document with particular attention paid to the role played by Unitrends' Secure Data Sync software that optimizes the transmission of data to the vault.

### INTRODUCTION

Natural disasters. Fire. Floods. Power failures. Theft. Viruses. Equipment failure. Malicious attacks. Human error. The list of reasons that disaster recovery is needed by any business seems endless.

The statistics are beyond sobering; for any business owner, they are terrifying. Of companies that had a major loss of computerized records, 6% survive long term.<sup>1</sup> Let that sink in a moment - 6%. The house edge in Vegas for roulette is only 2.7%. Of course, the "good" news is that 51% of these businesses reopen - but close within two years. 43% don't reopen - ever.

After the "blinding flash of the obvious" realization that one's business is at serious risk without a disaster recovery strategy the problem then shifts to the tactical issues of implementing a disaster recovery solution. The focus of this technology brief is to help the reader to gain increasing understanding of the practical considerations of implementing a Disaster Recovery solution with respect to the performance of Unitrends' Vaulting product in association with common long-haul networking technologies.

### DISASTER RECOVERY AND VAULTING

Unitrends provides a unique solution set for disaster recovery. Unlike software-only solutions, or service-only solutions, Unitrends employs a grid architecture that enables it to flexibly handle any number of networking topologies. The key software technology that makes this possible is known as Unitrends Secure Data Sync. This technology employs scheduling, throttling, compression, and de-duplication in order to optimize the amount of data that is required to be sent from an on-premises Data Protection Unit to an off-premises Data Protection Vault.

Unitrends offers through its resellers both the ability for a customer to purchase one or more Data Protection Units and one or more Data Protection Vaults or for the customer to purchase one or more Data Protection Units and use vaulting as a service through its resellers. In this topology, the Data Protection Units act as clients to the services oriented vault in a traditional client/server services architecture.

In addition, Unitrends allows "cross-vaulting" - a technique by which the on-premises Data Protection Units may assume a dual personality by adding a vaulting option. Thus, in the figure that follows, each of the five premises could be cross-vaulting to the other such that disaster recovery is achieved - in essence, rather than having a dedicated geographically non-proximate location dedicated solely to vaulting, a company can make use of its locations to lower the associated costs of the data center dedicated to vaulting. All of this is made possible by the Unitrends Secure Data Sync.

---

<sup>1</sup> Jim Hoffer, "Backing Up Business - Industry Trend or Event", Health Management Technology, January 2001.

# Unitrends™

In the diagram below, vaults may exist at each of the remote branches as well as the central premises with each branch cross-vaulting with the central premise (and vice versa.) Note that this is an arbitrary network configuration; each remote branch could link to and thus cross-branch with another remote branch as needed on a one-to-one basis. Unitrends Secure Data Sync allows an almost infinite flexibility regarding the provision of vaulting services rather than forcing customers into a hierarchical network model.



## FACTORS GOVERNING DISASTER RECOVERY

The three primary factors governing disaster recovery implementation with respect to performance are as follow:

- Network characteristics between the site that is being protected and the site which is the targeted repository of data of the protected site.
- The amount of data that is specified to be protected and the amount of data that periodically changes that is marked for disaster recovery.
- The software techniques used to effectively manage the bandwidth that is marked for disaster recovery (Unitrends Secure Data Sync.)

These factors are closely coupled; while each may be optimized to a limited degree individually the greatest impact on overall vaulting performance will be achieved by examining each factor in light of the other factors.

Each of these factors will be discussed in detail in the following chapters.



## NETWORK CHARACTERISTICS

Networks that are used for disaster recovery are typically more expensive, less reliable, and have lower bandwidth and higher latency than local area network links. Therefore it is important that the network type that best fits the requirements be selected, the amount of protected data be optimized, and that the software techniques used to transmit the data be efficient in order to decrease networking costs, increase networking reliability, optimize the use of available bandwidth, and handle wide variances with respect to latencies. This chapter is dedicated to discussing types of networks and their characteristics; the amount of data to be protected and the software techniques used to transmit the data are topics that are discussed in subsequent chapters.

## NETWORK TYPES AND COSTS

As noted previously, networks that are used for disaster recovery are typically more expensive than local area networks. This is exacerbated by increasing bandwidth requirements. The figure to the right depicts the typical relationship between the types of links most often used for disaster recovery and their associated pricing:<sup>2</sup>

Monthly Fractional & Full Internet Port Pricing		
Bandwidth	Monthly	Yearly
T1 (1.544 Mbps)	\$350	\$4,200
NxT1 (3 Mbps)	\$999	\$11,988
NxT1 (4.5 Mbps)	\$1,499	\$17,988
NxT1 (6 Mbps)	\$1,798	\$21,576
NxT1 (7.5 Mbps)	\$1,998	\$23,976
T3 (6 Mbps)	\$1,798	\$21,576
T3 (9 Mbps)	\$2,398	\$28,776
T3 (12 Mbps)	\$2,797	\$33,564
T3 (20 Mbps)	\$3,996	\$47,952
T3 (45 Mbps)	\$3,250	\$39,000
OC3 (155 Mbps)	\$18,681	\$224,172
OC12 (622 Mbps)	Call For Quote	
OC48 (2488 Mbps)	Call For Quote	
The above pricing is for the monthly port charges.		
Requires 1 Year Term.		

This seems very expensive...until you ask yourself what your downtime costs are per hour. IDC estimates that a typical medium-sized business has downtime costs that average \$78,000 per hour and that the typical medium-sized business loses more than \$1 million dollars per year due to downtime. And of course, as we saw above, a business that suffers a major loss of computerized records has a 6% chance of long-term survival.

Clearly, a smart business owner or IT leader is going to choose to implement a disaster recovery plan - but is going to optimize the return on investment such that the most important data is protected and that it is protected in a manner that optimizes network costs.

There are several other issues associated with the network types:

- A 1.544 MbPS T1 service means that all 24 individual channels constituting the T1 service are dedicated for transmission. Quite often, all 24 individual channels are not dedicated for transmission. In this case, the theoretical transport speed of the T1 drops accordingly.
- A bonded T1 service in which multiple T1 services are being provided does not necessarily mean that the T1 services will provide aggregate bandwidth consistent with the number of T1 services that are bonded. For example, in a bonded pair T1 service, we find that quite often one T1 service will be dedicated to inbound transmission and a second will be dedicated to outbound transmission. While this means that the odds of achieving T1-rated speeds for vaulting are better, it does not mean that you can exceed the single rated T1 transport bandwidth.
- When most technical people hear the term "T3" they believe that it automatically means a 45MbPS transport speed. T3 connections do not automatically provide 45MbPS of transport speed; there are different flavors ranging from 6 MbPS to 45 MbPS. On the positive side, T3 service does not suffer from the inbound/outbound separation issues that bonded T1 can potentially suffer.

<sup>2</sup> This pricing is for illustration purposes only; please check with your networking provider for actual costs.

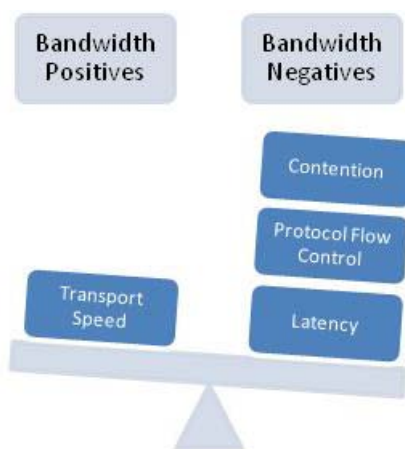
- OCx networks are common carrier fiber-based services and do not suffer from the inbound/outbound separation issues that bonded T1 can potentially suffer.

In addition, customers are increasingly using broadband services to replace T1, T3, and OCx connections. Broadband services are typically much less expensive, and the price/performance curve with respect to broadband services is climbing much faster than T1, T3, and OCx. However, broadband connections are much less predictable, and at times reliable, than T1, T3, and OCx services. In addition, broadband connections are typically asymmetric in terms of inbound/outbound performance with much higher inbound performance than outbound performance. This means that the data transmission from the premises to the “network cloud” (i.e., from the Data Protection Unit to the network which will eventually transport the data to the vaulting service) is much slower than the transmission of data in the opposite direction. In addition, broadband services may range quite a bit in effective transmission speed based on what other customers using that network service are transmitting. If a broadband service is being used as the network for vaulting, then special care should be taken to fully characterize the performance of the connection over a range of times and days.

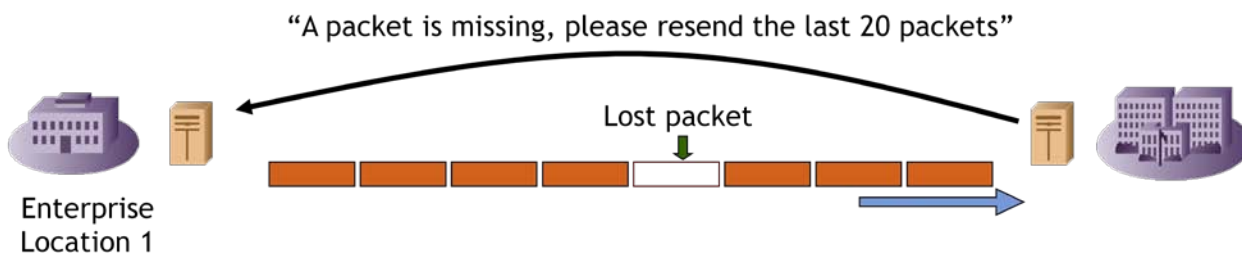
## NETWORK PERFORMANCE

The primary factors governing networking performance are

- Contention: When a number of applications are attempting to use the network at the same time this results in contention for the bandwidth.
- Latency: The delay introduced by the physical transport, the distance, and the intermediate equipment which decreases the response time of the applications using the network.
- Protocol Flow Control: The techniques used by the protocols that use the physical transport in order to optimize quality and performance of the data transmission.
- Transport Speed: The theoretical amount of data that a networking transport can be transmitted in a given time period.



Contention, latency, and protocol flow control are the enemies of the theoretical transport speed (see the figure to the left of this text). Contention is a particularly vexing problem with many protocols because the overhead associated with resolving contention is so much higher than you might think due to the design of these protocols. If we take TCP/IP flow control as an example, TCP enables guaranteed services across lossy IP networks and is essential for highly reliable data transmission. However, the way that TCP was designed will cause a rebroadcast message to be sent if a collision between two packets of information occurs due to contention (or a transport quality issue, etc.) TCP has a mechanism to request the retransmission of lost packets. The bandwidth wasted due to this mechanism increases as the data rate increases and the latency increases. This is because it takes time to recognize that a packet has been lost, time for the “lost packet” message to be sent, and time for the lost data to be resent. This is depicted in the figure below:



In multi-tenancy architectures, the effective bandwidth of the network tends to drop due to an increased possibility of “collisions” of packets over an incoming network link thus increasing protocol flow control operations.

What does all of this mean? It means that the actual bandwidth that may be achieved over a specified network will range from almost but not quite what is specified to much lower than is specified. The rule of thumb used to be that in order to compute effective bandwidth, you should assume a byte was 10-bits long and calculate the effective bandwidth on that basis.

If only life were that simple. It’s not a bad rule of thumb as rules of thumbs go; but it is hopelessly inadequate for real planning. This assumption of 25% doesn’t take into account heavily loaded networks, networks suffering quality issues, proximate latency issues, and so on. In the real world, you need to “measure twice and cut once” – in other words, you need to use measurement equipment to understand network performance and even then budget excess capacity for unintended events. For those unwilling to measure, or who have significantly varying network characteristics over time, the other rule of thumb is the 50% rule – the rule that most of the time a network can support 50% of its specified bandwidth.

Specified Bandwidth (MbPS)	Specified Bandwidth (MBPS)	“10-Bit Byte” Bandwidth (MBPS)	“10-Bit Byte” 100GB Time (Hours)	“50%” Bandwidth (MBPS)	“50%” 100GB Time (Hours)
1.5	0.19	0.15	185.19	0.09	296.30
3	0.38	0.30	92.60	0.19	148.15
6	0.75	0.60	46.30	0.38	74.07
10	1.25	1.00	27.78	0.63	44.44
45	5.63	4.50	6.17	2.81	9.88
100	12.50	10.00	2.78	6.25	4.44
155	19.38	15.50	1.79	9.69	2.87

*(Important note: This table represents the maximum capacity of the network and does not include any Secure Data Sync processing time. Secure Data Sync processing time will be discussed in the next chapter.)*

Note that the specified bandwidth is given in something called MbPS (which may also be referred to as “mbps”) which stands for megabits-per-second. We’ve then translated that into MBPS (megabytes-per-second) by dividing by 8 since there are 8 bits in byte. Seems simple, right? We’ve seen more confusion by people believing that an MbPS rating is a MBPS rating than most people would believe. In any case, conceptually, this isn’t a difficult concept.

The next two columns illustrate the consequences of the 10-bit byte rule discussed previously while the last two columns illustrate the consequences of the 50% bandwidth rule.

The formula for the 10-bit byte 100GB time calculation is given by

$$Time = \frac{Specified\ Bandwidth\ (Mbps)}{10 \times 60 \times 60}$$

The formula for the 50% 100GB time calculation is given by

$$Time = \frac{Specified\ Bandwidth\ (Mbps)}{8 \times 60 \times 60} \div 2$$

In our experience, real-life data transfer rates range between the 10-bit byte model and the 50% model; but your mileage will definitely vary.

What isn't in doubt, however, is that trying to vault 100GB of data across a 1.5MbPS network exceeds most backup windows (over 185 hours for the 10-bit byte model and over 230 hours for the 50% model). Fortunately, there are two techniques to radically reduce the amount of data that must be transferred on a daily basis: updating only the files that have changed and updating only the parts of the files that have changed. These are discussed in the next two chapters.

## PROTECTED DATA

Production data has increased at a 50% compounded annual growth rate over the last decade fueled by shrinking storage costs while non-proximate network bandwidth has seen relatively little improvement. The result is that there is a tremendous and growing amount of data that is subject to potential disaster recovery protection but there is a limited amount of bandwidth that may be used to protect that data.

If **all** production data changes on a daily basis, then it is much more cost-effective to backup to removable media and physically transport that media each day. Unitrends actually has options for this situation involving near line archiving at the Data Protection Unit level; however, it is incredibly rare for that situation to occur. Typically, it is much more cost-effective to use vaulting to automatically provide disaster recovery services.

Instead, what we want to do is to "seed" the vault with a master and then update the incremental changes on a periodic (typically daily) basis. Seeding may be done over a weekend if the network has enough bandwidth and the data to be protected is small enough; typically however, seeding is done by physically creating the master on the Data Protection Unit and then physically transporting that master to the vaulting service. Once seeding has occurred, only incremental changes to files need to be updated.

So what is the typical incremental change of data on a day-to-day basis? We find that typically the percentage of change is 2% to 4%; however, there are environments in which this is less and environments in which this is much more. The only way to really understand this is to actually measure the number and size of the files that have changed for several days.

Unitrends allows its customers to configure vaulting on a per-client level (a customer's server is a client of the Unitrends' Data Protection Unit) as well as on an entire Data Protection Unit basis. Thus, it is possible to designate several servers as being backed up by the Data Protection Unit but a subset of those servers as being vaulted. Often customers will designate their highest priority data on just a few servers that are then marked for vaulting as a way to manage the amount of data to be transmitted to the vault.

The following table depicts the effect of various amounts of incremental data change and the associated transmission times using as a basis the bandwidth table from the previous section and assuming 1 terabyte of production data.

Specified Bandwidth (Mbps)	Specified Bandwidth (MBPS)	"10-Bit Byte" Bandwidth (MBPS)	"10-Bit Byte" 100GB Time (Hours)	"50%" Bandwidth (MBPS)	"50%" 100GB Time (Hours)
1.5	0.19	0.15	185.19	0.09	296.30
3	0.38	0.30	92.60	0.19	148.15
6	0.75	0.60	46.30	0.38	74.07
10	1.25	1.00	27.78	0.63	44.44
45	5.63	4.50	6.17	2.81	9.88
100	12.50	10.00	2.78	6.25	4.44
155	19.38	15.50	1.79	9.69	2.87

(Important note: This table represents the maximum capacity of the network and does not include any Secure Data Sync processing time. Secure Data Sync processing time will be discussed in the next chapter.)

Clearly, the assumed amount of incremental change and the assumed effective bandwidth of the network are incredibly important in calculating the amount of network-related transmission time that will occur.

## UNITRENDS SECURE DATA SYNC

So, we have a terabyte of information we want to protect using a relatively efficient T1 line and we've measured and found that we have a 2% incremental data change. As we can see with the last table of the previous chapter, we are still looking at over 37 hours to transmit the changed data even if we assume that the T1 line is relatively efficient. Does this mean that we are then forced to spend the money to upgrade to a higher speed T1 or even a T3 or OC3?

Not necessarily - because Unitrends offers Secure Data Sync. Secure Data Sync is a technology that goes beyond file-level changes - it looks for block-level changes within files and even between files. Thus, it is possible to reduce the "incremental data change" shown in the previous chapter to a much smaller number based on the actual number of blocks that have changed.

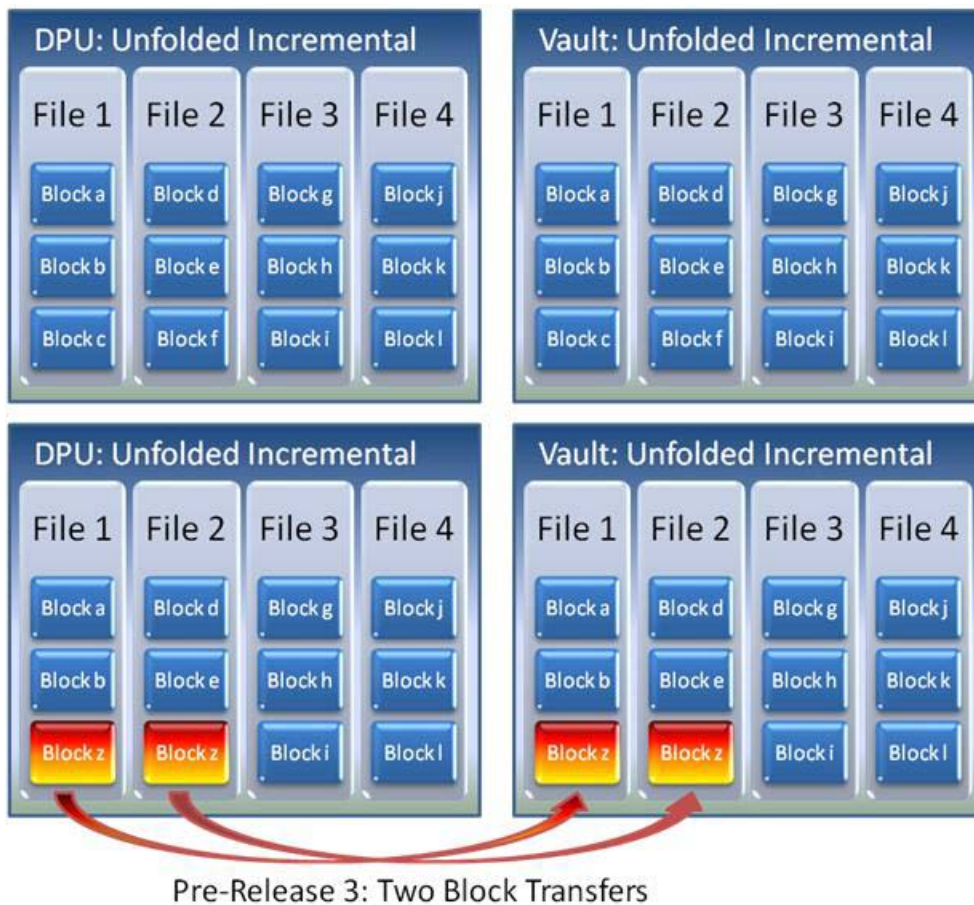
How does this help? Imagine that you have a large Microsoft Outlook ".PST" file - not exactly an unusual phenomenon.<sup>3</sup> If I add a single contact, or meeting, or message to my PST file that would mean that I just changed a file - and thus that entire file would be marked as having incrementally changed. That's expensive in terms of network performance.

However, in actuality only a very small number of bytes changed - the data I added plus any metadata that the Microsoft Outlook application uses to indicate the new data. If I can isolate that data and only have it vaulted instead of the entire PST file I've just reduced the amount of network bandwidth needed by a great margin. That's what Secure Data Sync does.

Of course, this begs a fundamental question: how large is a block? After all, a file could be considered just one large block, right? Well, the good news and the bad news is that there's no easy answer to this. That's bad news if you're writing a document describing Secure Data Sync; it is good news if you're actually using Secure Data Sync. That's because Secure Data Sync incorporates a technology known as "adaptive variable block sizing" that analyzes the data to be transmitted and selects an optimal block size based on that data. Block sizes range from 4KB to 2MB.

The figure below depicts an incremental backup being vaulted before release 3.

<sup>3</sup> I went and looked to find the size of my PST file while I was writing this document; it was slightly over 938MB and growing fast!



## RELEASE 3: THE NEXT STEP IN REDUCING NETWORK BANDWIDTH<sup>4</sup>

In previous releases, Unitrends' Secure Data Sync would perform its block-level examination on a file-by-file basis. This had the following disadvantages:

- A great deal of time was spent "unfolding" the master and incremental backups for a client so that file-by-file examination could be performed.
- A great deal of temporary storage was used unfolding the master and incremental backups.
- Only block matching within the same file could be de-duplicated.<sup>5</sup> This meant that blocks that matched between users (as a trivial example, several users with their own PST files who had the same large PowerPoint attachment) would still be transferred. Thus, as depicted in the preceding figure, if the same block exists in two different files then that block must be transferred twice.
- There was a tendency with respect to large files for the established link created over the network between the Data Protection Unit and the vaulting service to be dropped because all network activity would cease as

<sup>4</sup> This chapter applies to release 3 and all subsequent releases.

<sup>5</sup> Data de-duplication is all the rage; Unitrends has many competitors that go on and on about various forms of data de-duplication that they perform. At its heart, data de-duplication is a relatively simple operation - you simply look for blocks of information that are redundant and you don't copy those that are redundant. Unitrends has been doing data de-duplication for years and we are on our third generation of the technology.

the Data Protection Unit created the delta. This drop meant that all previous files that had been transmitted to the vault would have to be re-transmitted; basically, it was necessary to restart beginning with the first file of the backup. All of the files that had been previously transferred would be “forgotten” and this would have a tendency to use all available disk space of the vaulting service.

- Secure Data Sync made poor use of the available network capacity as it switched from processor-intensive delta creation on a per-file basis to the network-intensive task of transmitting those deltas. Basically, at times Secure Data Sync acted like a demanding child: it randomly wanted all of your attention or none of it but nothing in-between.

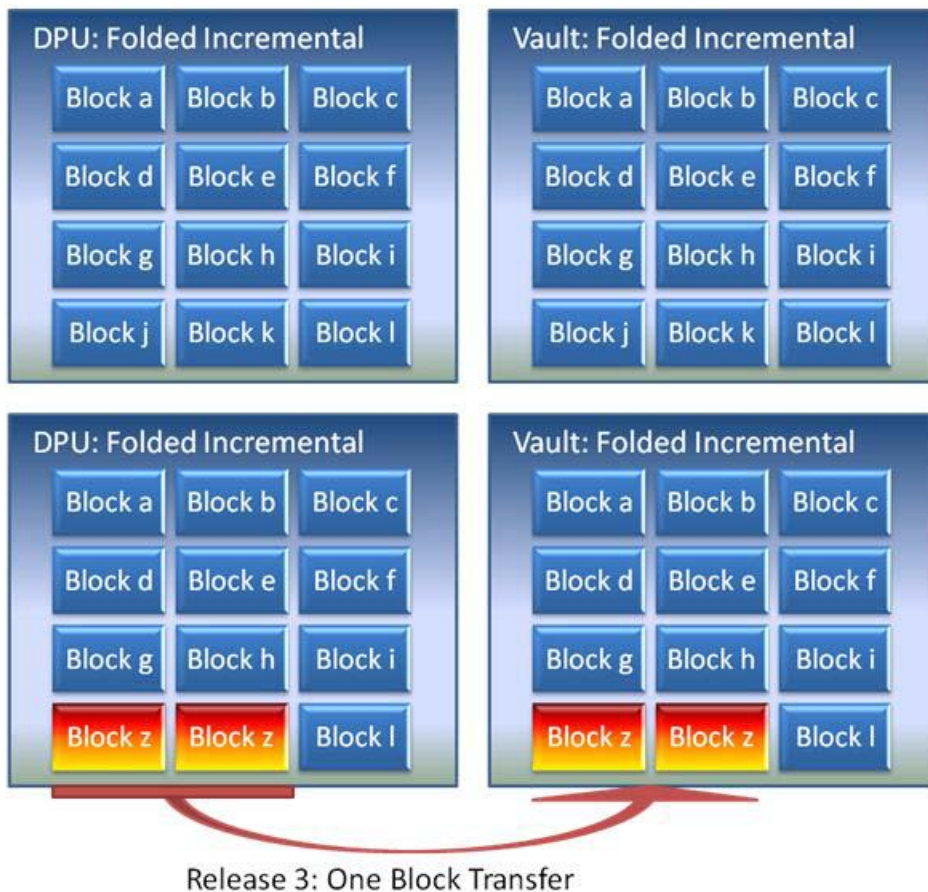
As a result, Unitrends created Release 3 of its Secure Data Sync technology. This technology had the following advantages over earlier versions:

- No time is spent unfolding or re-folding the master and incremental backups.
- No temporary storage is used for unfolding of the master and incremental backups.
- Data de-duplication (i.e., block matching) takes place across an entire client (i.e., the customer’s server) instead of across only a single file. Thus, as depicted in the figure following this text, if the same block exists in multiple files spanning a single server then that block is sent to the vault only once.
- Network usage is much more efficient since sporadic network usage no longer occurs and because less information must be transmitted due to the previously described client-wide data de-duplication.
- Large files no longer cause network links to be dropped.

What are the primary disadvantages of Release 3?

- For very high bandwidth networks (1Gbps and higher) that do not suffer from contention, the delta creation time can exceed the time it takes to simply copy the data over without de-duplication. Unitrends has several options for this eventuality including the elimination of de-duplication; for more information please see your Unitrends representative.
- There are relatively rare cases in which files are modified after a master but not modified during subsequent backup periods after that master and that the savings associated with backup-wide data de-duplication are offset by those files not having to be checked (not transmitted, but checked) previous to release 3. This has actually been modeled and found to have a very low incidence of occurrence both with our sample data sets and with industry-wide standard data access/modification patterns.
- It is disquieting for those used to previous releases to see all of the delta creation time to be spent up-front rather than having the delta creation time iterated sequentially with the data transmission (e.g., the trickle effect.)

The figure below depicts an incremental backup being vaulted in release 3.



## WHAT'S THE BOTTOM LINE?

No one can predict how much of a customer's replicated data is replicated blocks versus non-replicated blocks. Even more than understanding incremental data change, the replication within a data set depends completely on the actual contents of the customer's data. We recommend starting conservatively by vaulting data assuming relatively little data de-duplication and then increasing the amount of data being vaulted over time as experience and confidence with respect to the format of the customer's data increases.

## SUMMARY

Understanding vaulting performance is complicated because there are so many factors that have an influence on the time it takes to transmit changed data to the vault. As has been discussed in this paper, the three primary factors are the characteristics of the network used to transmit the data to the remote location, the amount of data that is specified to be protected and how much that data periodically changes and the software techniques used to effectively manage the bandwidth that is marked for disaster recovery.