

NON-LOCAL REDUNDANCY

ERASURE CODING AND DISPERSED REPLICAS
FOR ROBUST RETRIEVAL
IN THE SWARM PEER-TO-PEER NETWORK

Swarm Research Division^{*†‡}

v1.0 – Jan 2024



Contents

1	Error correcting codes	2
2	Erasures coding in the Swarm hash tree	3
3	Levels of security and the number of parities	4
4	Dispersed replicas	9
5	Prefetching strategies for retrieval	12

*Viktor Trón, György Barabás, Viktor Tóth, Callum Toner, Dan Nickless

†Two former members of Swarm Research, Dániel A. Nagy and Áron Fischer, have made substantial contributions to the solution presented here.

‡The Swarm authors are indebted to Andrea Robert whose thorough review greatly improved the paper.

First, in Section 1, we introduce erasure codes, and show how to apply them to files in Swarm (Section 2). In Section 3, we introduce security levels of data availability and derive their respective parameterisations. In Section 4, we describe a construct that enables cross-neighbourhood redundancy for singleton chunks and which completes erasure coding. Finally, in 5, we propose a number of retrieval strategies applicable to erasure-coded files.

1 Error correcting codes

Error correcting codes are widely utilised in the context of data storage and transfer to ensure data integrity even in the face of a system fault. Error correction schemes define how to rearrange the original data by adding redundancy to its representation before upload or transmission (*encoding*) so that it can correct corrupted data or recover missing content upon retrieval or reception (*decoding*). The different schemes are evaluated by quantifying their strength (*tolerance*, in terms of the rate of data corruption and loss) as a function of their cost (*overhead*, in terms of storage and computation).

In the context of computer hardware architecture, synchronising arrays of disks is crucial for providing resilient storage in data centres. In *erasure coding*, in particular, the problem can be framed as follows: How does one encode the stored data into shards distributed across the disks so that the data remains fully recoverable in the face of an arbitrary probability that any one disk becomes faulty? Similarly, in the context of Swarm’s distributed immutable chunk store, the problem can be reformulated as follows: How does one encode the stored data into chunks distributed across neighbourhoods in the network so that the data remain fully recoverable in the face of an arbitrary probability that any one chunk is not retrievable?¹

Reed-Solomon coding, or RS coding for short (Reed and Solomon 1960, Bloemer et al. 1995, Plank and Xu 2006, Li and Li 2013), is the father of all error correcting codes and is also the most widely used and implemented one.² When applied to data of m fixed-size blocks (message of length m), it produces an encoding of $m + k$ *codewords* (blocks of the same size) in such a way that having any m out of $m + k$ blocks is sufficient to reconstruct the original data. Conversely, k puts an upper bound on the number of *erasures* allowed (number of blocks unavailable) for full recoverability, i.e., it expresses (the maximum) *loss tolerance*.³ k is also the count of *parities*, quantifying the data blocks added during the encoding on top of the original volume, i.e., it expresses *storage overhead*. While RS is, therefore, optimal for storage (since loss tolerance cannot exceed the storage overhead), it has high bandwidth demands⁴ for local repair processes.⁵ The decoder needs to retrieve m chunks to recover a particular unavailable chunk. Hence, ideally, RS is used on files which are supposed to be downloaded in full,⁶ but it is inappropriate for use cases needing only local repairs.⁷

When using RS, one common option is to use *systematic* encoding, which means that the original data form part of the encoding. That is, the parities are actually added to them. This is the option we use in Swarm.

¹We will assume that the retrieval of any one chunk fails with equal and independent probability.

²For a thorough comparison of an earlier generation of implementations of RS, see Plank et al. (2009).

³Error correcting codes that have a focus on correcting data loss are referred to as *erasure codes*, a typical scheme of choice for distributed storage systems (Balaji et al. 2018).

⁴Both the encoding and the decoding of RS codes takes $O(mk)$ time (with m data chunks and k parities). However, we found computational overhead both insignificant for a network setting as well as undifferentiating.

⁵Entanglement codes (Estrada-Galinas et al. 2018, 2019) require a minimal bandwidth overhead for a local repair, but at the cost of storage overhead that is in multiples of 100%.

⁶Or in fragments large enough to include the data span over which the encoding is defined, such as videos.

⁷Use cases requiring random access to small amounts of data (e.g., path lookup) benefit from simple replication to optimise on bandwidth, which is suboptimal in terms of storage (Weatherspoon and Kubiatowicz 2002).

2 Erasure coding in the Swarm hash tree

Swarm uses the *Swarm hash tree* to represent files. This structure is a Merkle tree (Merkle 1980), whose leaves are the consecutive segments of the input data stream. These segments are turned into chunks and are distributed among the Swarm nodes for storage. The consecutive chunk references (either in the form of an address or an address and an encryption key) are written into a chunk on a higher level. These so-called *packed address chunks* (PACs) constitute the intermediate chunks of the tree. The branching factor b is chosen so that the references to its children fill up a full chunk. With a reference size of 32 or 64 (hash size 32) and a chunk size of 4096 bytes, the value of b is 128 for plain, and 64 for encrypted content (Figure 1).

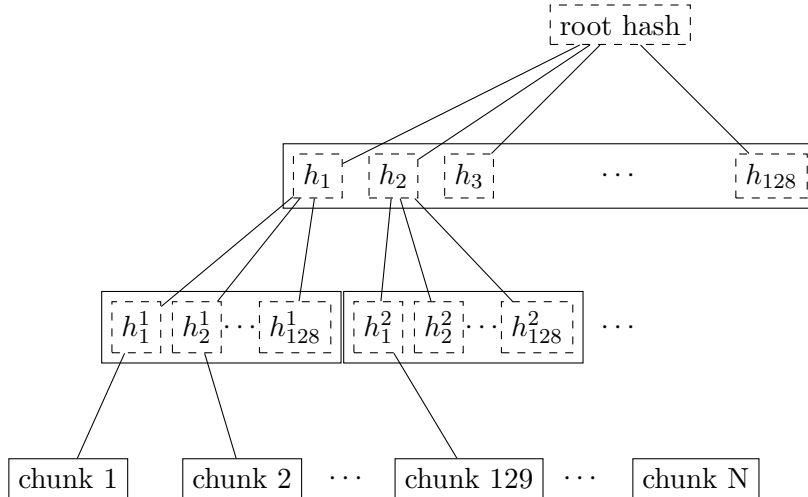


Figure 1: The Swarm tree is the data structure encoding how a document is split into chunks.

Note that on the right edge of the hash tree, the last chunk of each level may be shorter than 4K: in fact, unless the file is exactly $4 \cdot b^n$ kilobytes long, there is always at least one *incomplete chunk*. Importantly, it makes no sense to wrap a single chunk reference in a PAC, so it is attached to the first level where there are open chunks. Such *"dangling" chunks* will appear if and only if the file has a zero digit in the b -ary representation of its length in chunks.

During file retrieval, a Swarm client starts from the root hash reference and retrieves the corresponding chunk. Interpreting the metadata as encoding the span of data subsumed under the chunk, it decides that the chunk is a PAC if the span exceeds the maximum chunk size. In case of standard file download, all the references packed within the PAC are followed, i.e., the referenced chunk data are retrieved.

PACs offer a natural and elegant way to achieve consistent redundancy within the Swarm hash tree. The input data for an instance of erasure coding is the chunk data of the children, with the equal-sized bins corresponding to the chunk data of the consecutive references packed into it. The idea is that instead of having each of the b packed references represent children, only m do, and the rest of the $k = b - m$ encode RS parities (Figure 2).

The *chunker* algorithm that incorporates PAC-scoped RS encoding works as follows:

1. Set the input to the actual data level and produce a sequence of chunks from the consecutive 4K segments of the data stream. Choose m and k such that $m + k = b$ is the branching factor (128 for plain, and 64 for encrypted content).
2. Read the input one chunk at a time. Count the chunks by incrementing a counter i .
3. Repeat Step 2 until either $i = m$ or there are no more data left.
4. Use the RS scheme on the last $i \leq m$ chunks to produce k parity chunks resulting in a total of $n = i + k \leq b$ chunks.
5. Concatenate the references of all these chunks to result in a packed address chunk (of size $h \cdot n$) on the level above. If this is the first chunk on that level,— set the input to this level and spawn this same procedure from Step 2.
6. When the input is consumed, signal the end of input to the next level and quit the routine. If there is no next level, record the single chunk as the root chunk and use the reference to refer to the entire file.

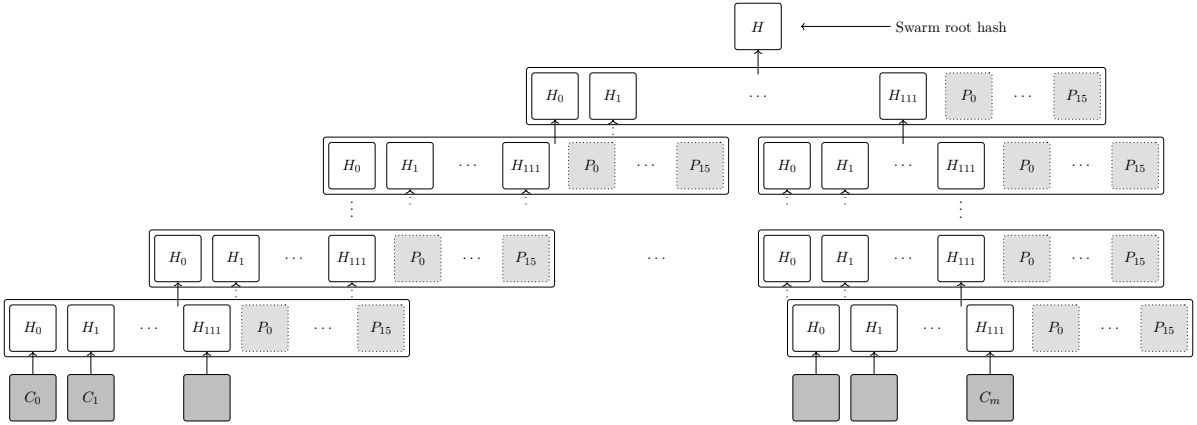


Figure 2: The Swarm tree with extra parity chunks using $m = 112$ out of $n = 128$ RS encoding. Chunks P_0 through P_{15} are parity data for chunks H_0 through H_{111} on every level of intermediate chunks.

This pattern repeats itself all the way down the tree. Thus, hashes H_{m+1} through H_{127} point to parity data for chunks pointed to by H_0 through H_m .⁸

3 Levels of security and the number of parities

Non-local redundancy is presented here as a scheme of encoding that allows strategies of retrieval in order to guarantee data availability. With packed address chunks set as the scope of erasure codes, it is crucial that we use the right number of shards and parities among the children of an intermediate node in the Swarm hash tree representing a file. Given assumptions about chunk retrieval error rates and the number of parities used, one can calculate the degree of certainty that the data can be recovered without error. One can even apply the same logic backwards: given some level of certainty with which we want recovery to be error-free, we can compute how many parities should be used to provide that level of safety. In what follows, we give a formal exposition of how to find these parity counts.

Let there be m original chunks and k parity chunks, such that any m chunks out of the total $n = m + k$ ones are fully recoverable after the loss of any k of them. In the process of retrieving

⁸Since parity chunks P_i do not have children, the tree structure does not have uniform depth.

the n chunks, what is the likelihood of overall data corruption, given a per-chunk probability of error ϵ ?

By “overall data corruption”, we mean that more than k chunks are damaged in the data retrieval process. We assume that each chunk’s probability of error is independent of other chunks. In that case, the problem boils down to the independent drawing of n chunks, each of which undergo a *Bernoulli trial* of being faulty with probability ϵ . The total number of faulty chunks out of n independent Bernoulli trials is given by the *binomial distribution*:

$$B(i, n, \epsilon) = \binom{n}{i} \epsilon^i (1 - \epsilon)^{n-i}. \quad (1)$$

This expression is the probability mass function for the binomial distribution, yielding the probability that out of n chunks, exactly i will be faulty—assuming that the per-chunk probability of error is ϵ .

Since there are k parities out of the n chunks, the system can tolerate up to k chunk errors. The probability that no more than k errors accumulate can be expressed by summing Equation 1 over i up to k :

$$P(k, n, \epsilon) = \sum_{i=0}^k \binom{n}{i} \epsilon^i (1 - \epsilon)^{n-i}, \quad (2)$$

which is the cumulative distribution function of the binomial distribution.

One typical question is the following: given the number of chunks n and a value α such that we want the overall probability of data corruption to be below this value, how many out of the n chunks should be parities? Since $P(k, n, \epsilon)$ is the probability that *no more than* k errors accumulate, $1 - P(k, n, \epsilon)$ is the probability of more than k errors; i.e., that *at least* $k + 1$ errors accumulated and therefore the data are corrupted. We want to keep this probability below α , so we can write

$$\alpha \geq 1 - P(k, n, \epsilon). \quad (3)$$

Rearranging, we have

$$1 - \alpha \leq P(k, n, \epsilon). \quad (4)$$

That is, we are looking for values of k which will satisfy this inequality (Figure 3). This can be obtained by inverting the cumulative distribution function in k , resulting in the quantile function $Q(1 - \alpha, n, \epsilon)$. While this inverse has no convenient closed-form expression, it can be efficiently evaluated numerically for any set of input parameters. As with any cumulative distribution function, $P(k, n, \epsilon)$ is monotonically increasing in k . Applying the inverse on both sides of Equation 4 therefore does not flip the direction of the inequality, and gives $k \geq Q(1 - \alpha, n, \epsilon)$. Or if we look for the smallest k satisfying this condition:

$$k = Q(1 - \alpha, n, \epsilon). \quad (5)$$

This is the formula yielding the minimum number of required parities to keep the overall probability of data corruption below α .

Figure 4 presents the number of parities needed as a function of error rate for various levels of security. Figure 5 presents the number of parities needed to keep the probability of overall data corruption at a given level for various values of the per-chunk error rate.




Figure 3: The point at $k = 17$ along the binomial distribution, where the probability of exceeding this many errors becomes less than $\alpha = 10\%$. Here, the total number of chunks is $n = 128$, and the per-chunk error rate is $\epsilon = 0.1$.




Figure 4: The number of parities needed (ordinate) as a function of the per-chunk error rate ϵ (abscissa), for keeping the probability of overall data corruption below given limits (colours) and for $n = 64$ chunks (left panel) and $n = 128$ chunks (right panel).

The same type of problem can also be phrased slightly differently: given a number of chunks n , how many parities k should be added to them to keep the overall data corruption probability below some level α ? In this case the total number of chunks is $n + k$ (instead of having n chunks, out of which k are parities), and so Equation 5 is modified to be

$$k = Q(1 - \alpha, n + k, \epsilon). \quad (6)$$

While this equation has no closed-form solution for k , one can easily find the k satisfying it as long as k is bounded in a relatively small range. In our case, the maximum number of chunks, $n + k$, is 128, and so k is at most $128 - n$. This makes it simple to find the value of k compatible with Equation 6. The number of parities in Tables 1-3 were obtained using this method.

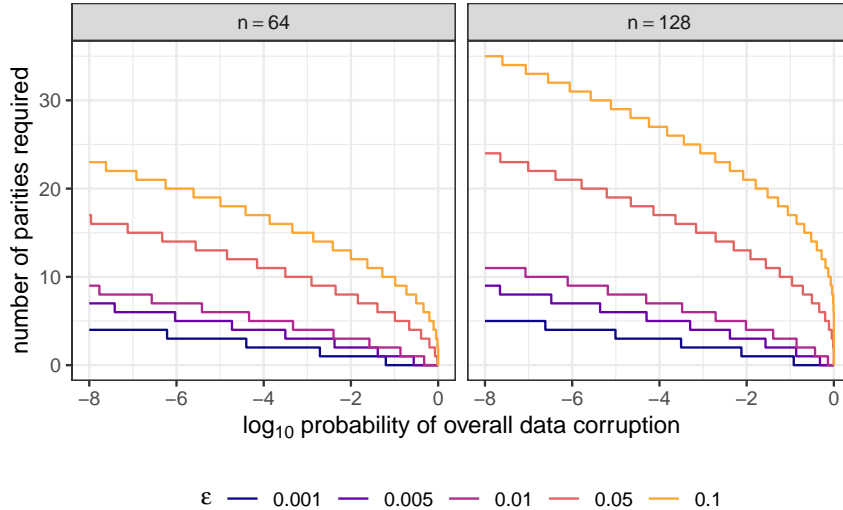


Figure 5: The number of parities required (ordinate) to keep the probability of overall data corruption at a given level (abscissa), for various values of the per-chunk error rate ϵ (colours) and for $n = 64$ chunks (left panel) and $n = 128$ chunks (right panel).

In principle, the exact parity counts can be made user-configurable. However, to make non-local redundancy a transparent and easy-to-use feature, we opted for a simplified yet intuitive interface. First of all, we set our maximum tolerated error rate of integrity at 10^{-6} , in other words our security constant expressing our certainty at 6 nines, 99.9999%. Second, we propose to use a handful of named security levels of (non-local) redundancy which correspond to assumptions about the maximum error rates of individual chunk retrievals expressed as discrete percentages. Table 1 lists the security levels with the corresponding assumption about the maximum error rate of chunk retrieval.

security		error rate of chunk retrieval	plain		encrypted	
level	name		chunks	parities	chunks	parities
0	NONE	0%	128	0	64	0
1	MEDIUM	1%	119	9	59	9
2	STRONG	5%	107	21	53	21
3	INSANE	10%	97	31	48	31
4	PARANOID	50%	38	90	19	90

Table 1: Security levels for non-local redundancy UI and corresponding assumptions about uniform and independent error rates of individual chunk retrieval. In subsequent columns we specify the composition of full chunks for the security levels for plain (columns 4 and 5) and encrypted (columns 6 and 7) content.

If the number of file chunks is not a multiple of m , it is not possible to proceed with the last batch in the same way as the others. We propose that we encode the remaining chunks with an erasure code that guarantees at least the same level of security as the others.⁹ Overcompensating, we could still require the same number of parity chunks even when there are fewer than m data

⁹Note that this is not as simple as choosing the same redundancy. For example, a 50-out-of-100 encoding is much more secure against loss than a 1-out-of-2 encoding, even though the redundancy is 100% in both cases.

chunks in the scope. However, we can also just calculate the necessary parities for all possible incomplete chunks and security levels. Figure 6 plots the number of parities against the number of chunks required:

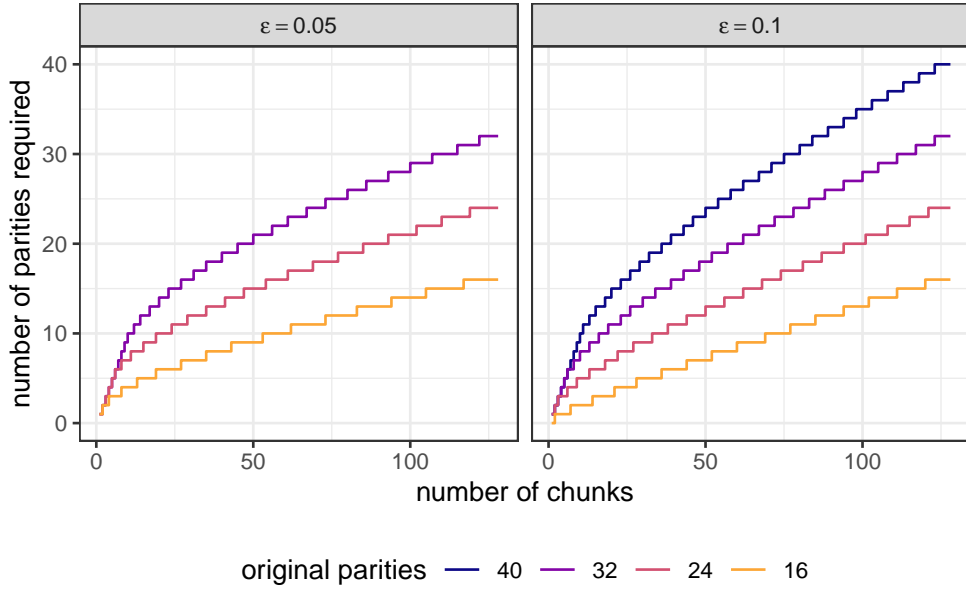


Figure 6: Number of chunks (abscissa) and the corresponding required number of parities (ordinate) such that will maintain the same overall probability of no data corruption as would be the case with 128 chunks, an original number of parities indicated by the colours, and a likelihood ϵ of an erroneous retrieval of a single chunk indicated in the panel headers.

Tables 2 and 3 show the number of chunks that are maintainable for a given number of parities k across various security levels. Since encrypted chunks are referenced with the hash address followed by the decryption key, an encrypted reference takes up 2 hash-sized segments. Parity chunks added to an encrypted PAC, however, are calculated based on the encrypted shards and are themselves not encrypted, hence their references only use a single hash. Thus, the number of effective hash-sized segments used is obtained as twice the number of chunks plus the number of parities. Since this can be an odd number and less than 128, in some security levels even the full chunks are not completely full.

As a final note, one should keep in mind that the probability of a failed data retrieval, $\alpha = 10^{-6}$, is not the same as the probability of a failed file retrieval. This is because α is only valid for one 128-chunk segment (64-chunk segment for encrypted content) of a file, not a file as a whole in general. Assuming that retrieval errors may occur independently to any chunk, we can use α and the size of a file to calculate the probability that a file as a whole is successfully retrieved. This probability is $1 - \alpha$ for each 128-chunk segment of a file, so if a file consists of s 128-chunk segments, then the probability is $(1 - \alpha)^s$. In terms of bytes: a file of g bytes consists of $g/2^{12}$ chunks (because 2^{12} bytes is 4KB), which then make up for $s = g/(2^{12} \cdot 2^7)$ 128-chunk segments (because $128 = 2^7$). This means that the probability P_F of a successful file retrieval is

$$P_F = (1 - \alpha)^{g/2^{19}}, \quad (7)$$

an exponentially decreasing function of the file size g . For example, a file of 1GB ($s = 2^{30}$ bytes) with $\alpha = 10^{-6}$ has $P_F = 0.998$, for a failure probability of $1 - P_F = 0.2\%$.

MEDIUM			INSANE		
parities	chunks		parities	chunks	
	plain	encrypted		plain	encrypted
2	1	-	5	1	-
3	2-5	1-2	6	2	1
4	6-14	3-7	7	3	1
5	15-28	7-14	8	4-5	2
6	29-46	14-23	9	6-8	3-4
7	47-68	23-34	10	9-10	4-5
8	69-94	34-47	11	11-13	5-6
9	95-119	47-59	12	14-16	7-8
			13	17-19	8-9
			14	20-22	10-11
			15	23-26	11-13
			16	27-29	13-14
			17	30-33	15-16
			18	34-37	17-18
			19	38-41	19-20
			20	42-45	21-22
			21	46-50	23-25
			22	51-54	25-27
			23	55-59	27-29
			24	60-63	30-31
			25	64-68	32-34
			26	69-73	34-36
			27	74-77	37-38
			28	78-82	39-41
			29	83-87	41-43
			30	88-92	44-46
			31	93-97	46-48

STRONG		
4	1	-
5	2-3	1
6	4-6	2-3
7	7-10	3-5
8	11-15	5-7
9	16-20	8-10
10	21-26	10-13
11	27-32	13-16
12	33-39	16-19
13	40-46	20-23
14	47-53	23-26
15	54-61	27-30
16	62-69	31-34
17	70-77	35-38
18	78-86	39-43
19	87-95	43-47
20	96-104	48-52
21	105-107	52-53

Table 2: The number of parities (first column in each table) to be appended to a given number of chunks (second and third column of each table, given as a range) so that the probability of an unsuccessful data retrieval remains below $\alpha = 10^{-6}$. The second column is for plain chunks, while the third one is for encrypted chunks. The tables are for security levels 1-3, to be continued for security level 4 in Table 3.

4 Dispersed replicas

This leaves us with only one corner case: it is not possible to use our m -out-of- n scheme on a single chunk ($m = 1$) because it would amount to $k + 1$ copies of the same chunk. The problem is that copies of the same chunk all have the same hash and therefore are automatically deduplicated. Whenever a single chunk is left over ($m = 1$) (i.e., the root chunk itself), we would need to replicate the chunk in a way that (1) ideally, the replicas are dispersed in the address space in a balanced way, yet (2) their address can be known by retrievers who ideally only know the reference to the original chunk's address.

PARANOID			PARANOID (continued)		
parities	chunks		parities	chunks	
	plain	encrypted		plain	encrypted
19	1	-	61	20	10
23	2	1	63	21	10
26	3	1	65	22	11
29	4	2	66	23	11
31	5	2	68	24	12
34	6	3	70	25	12
36	7	3	71	26	13
38	8	4	73	27	13
40	9	4	75	28	14
43	10	5	76	29	14
45	11	5	78	30	15
47	12	6	80	31	15
48	13	6	81	32	16
50	14	7	83	33	16
52	15	7	84	34	17
54	16	8	86	35	17
56	17	8	87	36	18
58	18	9	89	37	18
59	19	9	90	38	19

Table 3: As Table 2, but for the PARANOID security level.

Our solution uses Swarm’s special construct, the *single owner chunk* (SOC; Figure 7). Replicas of the root chunk are created by making the chunk data the payload of a number of SOCs. The addresses of these SOCs must be derivable from the original root hash following a deterministic convention shared by uploaders and downloaders.

The address of a SOC is the hash of its ID and the Ethereum address of its owner. In order to create valid SOCs, uploaders need to sign the SOC with the owner’s identity, therefore the owner of the SOC must be a consensual identity with their private key publicly revealed.¹⁰

The other component of the address, the SOC ID, must satisfy two criteria: (1) it needs to match the payload hash up to 31 bytes and (2) it must provide the entropy needed to mine the overall chunk into a sufficient number of distinct neighbourhoods. (1) is added as a validation criterion for the special case of replica SOCs, while (2) takes care that we can find replicas uniformly dispersed within the address space. This construct is called *dispersed replica*:

Let us assume c is the content-addressed chunk we need to replicate; n is the number of bits of entropy available to find the nonces that generate 2^k perfectly balanced replicas; initialise a chunk array ρ of length 2^k and start with n -bit integer $i = 0$ and replica counter $C = 0$.

1. Create the SOC ID by taking $addr(c)$ and changing the last byte (at index position 31) to i .

¹⁰This has the added benefit that third parties can also upload replicas of any chunk.

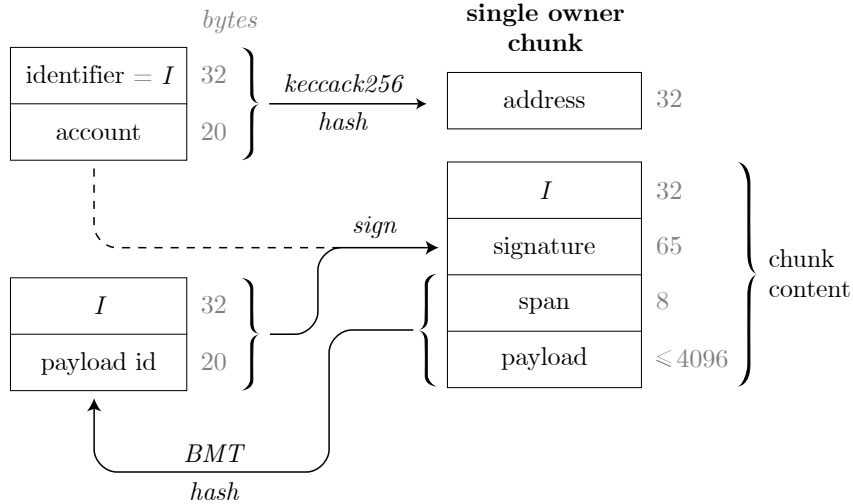


Figure 7: Single owner chunk (SOC). Unlike content-addressed chunks, SOCs obtain their integrity through the signature of their (single) owner and cross-owner immutability through hashing the owner’s address in the chunk address (effectively achieving access control via namespacing).

2. Calculate the the SOC address by concatenating ID id and owner o and hash the result using the Keccak256 base hash $a_i := H(id \oplus o)$, and record $c_i = SOC(id, o, c)$.
3. Calculate the bin this hash belongs to by taking the k -bit prefix as big-endian binary number j between $0 \leq j < 2^k$.
4. If $\rho[j]$ is unassigned, then let $\rho[j] := c_i$ and increment C .
5. If $C = 2^k$, then quit.
6. Increment i by one, if $i = 2^n$, then quit.
7. Repeat from Step 1.

With this solution, we are able to provide an arbitrary level of redundancy for the storage of data of any length.¹¹

Then, depending on the strategy, the downloader can choose which address to retrieve the chunk from. The obvious choice is the replica closest to the requesting node’s overlay address. In other words, the last item of the sorted chunk array ρ using the comparison function:

$$i < j \Leftrightarrow PO(Overlay(node), ADDRESS(\rho[i])) < PO(Overlay(node), ADDRESS(\rho[j])) \quad (8)$$

If the probability of any replica being faulty is ϵ , then, assuming independence, the probability that n parities are faulty is ϵ^n . Here we can write $n = k + 1$; that is, we have one “original” chunk and the rest of them are the k parities. Keeping the overall error probability below α then means that

$$\epsilon^{k+1} = \alpha \quad (9)$$

¹¹Note that if n is small, then generating all 2^k balanced replicas may not be achievable, and if $n < k$, this is certainly not possible. In general, given n, k at least m miss has a probability of $(1 - m/2^k)^{2^n}$.

must be satisfied. Taking logarithms on both sides and rearranging, we get

$$k = \frac{\log(\alpha)}{\log(\epsilon)} - 1. \quad (10)$$

This is the number of parities of a singleton chunk required to keep the overall data corruption probability below α . The base of the log in Equation 10 is arbitrary. This means that if we use base-10 logarithms and assume that $\alpha = 10^{-6}$, we get the simpler

$$k = \frac{6}{|\log_{10}(\epsilon)|} - 1. \quad (11)$$

For example, if the per-chunk error rate is ten percent ($\epsilon = 0.1$), then $|\log_{10}(\epsilon)| = |\log_{10}(1/10)| = 1$, and so $k = 6/1 - 1 = 5$ parities are needed. If instead the per-chunk error rate is just one percent ($\epsilon = 0.01$), then only $k = 6/2 - 1 = 2$ parities are necessary.

In particular, for the same per-chunk error rates as in Table 1, we get:

security level	error rate	parities required	dispersed replicas
NONE	0%	0	0
MEDIUM	1%	2	2
STRONG	5%	4	4
INSANE	10%	5	8
PARANOID	50%	19	16

Table 4: For a given per-chunk error rate (column), how many parities (second column) are required of a single chunk to keep the overall data corruption probability below $\alpha = 10^{-6}$?

5 Prefetching strategies for retrieval

When downloading, systematic per-level erasure codes allow for different *prefetching strategies*:

NONE = *direct with no recovery; frugal*

No prefetching takes place, RS parity chunks are ignored if present. Retrieval involves only the original chunks, no recovery.

DATA = *prefetching data but no recovery; cheap*

Prefetching data-only chunks, RS parity chunks are ignored if present, no recovery.

PROX = *distance-based selection; cheap*

For all intermediate chunks, first retrieve m chunks that are expected to be the fastest to download (e.g., the m closest to the node).

RACE = *latency optimised; expensive*

Initiate requests for all chunks within the scope ($\max m + k$) and will need to wait only for the first m chunks to be delivered in order to proceed. This is equivalent to saying that the k slowest chunk retrievals can be ignored, therefore this strategy is optimal for latency at the expense of cost.

All in all, strategies using recovery can effectively overcome the occasional unavailability of chunks, be it due to faults such as network contention, connectivity gaps in the Kademlia table, node churn, overpriced neighbourhoods, or even malicious attacks targeting a specific neighbourhood.

Similarly, given a typical model of network latencies for chunk retrieval, erasure codes in RACE mode can guarantee an upper limit on retrieval latencies.¹²

References

- Balaji, S., Krishnan, M.N., Vajha, M., Ramkumar, V., Sasidharan, B., Kumar, P.V., 2018. Erasure coding for distributed storage: An overview. *Science China Information Sciences* 61, 1–45.
- Bloemer, J., Kalfane, M., Karp, R., Karpinski, M., Luby, M., Zuckerman, D., 1995. An xor-based erasure-resilient coding scheme. Technical Report. International Computer Science Institute. Technical Report TR-95-048.
- Estrada-Galinanes, V., Miller, E., Felber, P., Pâris, J.F., 2018. Alpha entanglement codes: practical erasure codes to archive data in unreliable environments, in: 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), IEEE. pp. 183–194.
- Estrada-Galinanes, V., Nygaard, R., Tron, V., Saramago, R., Jehl, L., Meling, H., 2019. Building a disaster-resilient storage layer for next generation networks: The role of redundancy. IEICE Technical Report; IEICE Tech. Rep. 119, 53–58.
- Li, J., Li, B., 2013. Erasure coding for cloud storage systems: A survey. *Tsinghua Science and Technology* 18, 259–272.
- Merkle, R.C., 1980. Protocols for public key cryptosystems, in: Proc. 1980 Symposium on Security and Privacy, IEEE Computer Society, IEEE. p. 122.
- Plank, J.S., Luo, J., Schuman, C.D., Xu, L., Wilcox-O’Hearn, Z., et al., 2009. A performance evaluation and examination of open-source erasure coding libraries for storage., in: FAST, pp. 253–265.
- Plank, J.S., Xu, L., 2006. Optimizing Cauchy Reed-Solomon codes for fault-tolerant network storage applications, in: Network Computing and Applications, 2006. NCA 2006. Fifth IEEE International Symposium on, IEEE. pp. 173–180.
- Reed, I.S., Solomon, G., 1960. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics* 8, 300–304.
- Weatherspoon, H., Kubiatowicz, J.D., 2002. Erasure coding vs. replication: A quantitative comparison, in: Peer-to-Peer Systems: First International Workshop, IPTPS 2002 Cambridge, MA, USA, March 7–8, 2002 Revised Papers 1, Springer. pp. 328–337.

¹²For instance, in the temporally sensitive case of real-time video streaming, for any quality setting (bitrate and FPS), buffering times can be guaranteed if the batch of chunks representing a time unit of media is encoded using its own scope(s) of erasure coding.