Kryptos K4: The W-Rectangle

Read in Rows of 19, K4 Puts Four W's on a Rectangle and Its Cribs on the Edges

Guillaume Lethuillier
glthr.com

November 11, 2026 · updated November 28, 2026

Series: Kryptos K4 — Part 2 of 3
Next: Kryptos K4: Why the 'W' Matters (glthr.com/kryptos-k4-why-the-w-matters)
Interactive view: glthr.com/k4

Abstract

Written in rows of 19, the first 95 letters of Kryptos K4 put four of its five W's at the corners of a rectangle, with both of Jim Sanborn's clues lying along its edges. A half-turn of the grid sends 44 letters of one group from my 2024 article onto 44 letters of the other, leaving only the central ZZ and the final AR unexplained. None of this deciphers K4, and the rarity of the geometry does not prove it was intentional.

> **Update (November 28, 2026):** I revised this article: a shorter title, clearer explanations, a note on related earlier work, and Group B in teal, as in the rest of the series.

What happens if you copy the last message on the Kryptos sculpture and change only the line breaks? All 97 letters stay in order; what changes is which letters sit above one another.

In my 2024 article, I used the five W's to split K4 into six pieces. Assigned alternately to two groups, they form two halves of 46 letters whose letter counts match in shape: for every letter that appears three times in one group, for instance, another letter appears three times in the other. Members of the Kryptos group had already described this pattern in 2016.

Two years later, I came back to those five W's with a different question: what shape would they make on the page?

Written in rows of 19 letters, the first 95 letters of K4 form a grid of 5 rows and 19 columns. Four W's mark the corners of a rectangle; the fifth sits one cell from the grid's center. The two known plaintext clues, or *cribs*, lie on the edges of this rectangle. The same geometry sends 44 letters of Group A onto 44 letters of Group B under a half-turn (see the animation below).

Cite as: Guillaume Lethuillier, "Kryptos K4: The W-Rectangle", glthr.com, November 11, 2026. https://glthr.com/kryptos-k4-the-w-rectangle
Interactive figures, animations and videos of the web page are shown here as still images or addresses.

A Rectangle Hidden in K4

Let's write the first 95 letters of K4 in rows of 19. On the interactive page, you can watch the letters of K4 leave the copper for these rows.

On the copper, K4 begins at the end of one line and fills three more lines of 31 letters, as the first figure of the 2024 article shows. I found the rectangle by trying every width from 3 to 48.[1] Each width leaves its own tail outside the grid (2 letters at width 19, up to 31 at width 33), and the counts below take this into account.



- **Four W's form a rectangle.** They sit one cell in from each corner of the grid: at (row 2, column 2), (2, 18), (4, 2) and (4, 18).
- **The fifth W marks the middle.** At (3, 11), it is one cell right of the center cell of the grid. It is also the 49th of the 97 letters: the exact middle of K4.
- **The two cribs lie on the edges of the rectangle.** EASTNORTHEAST runs along the top edge and starts in the cell right after a W; BERLINCLOCK runs along the bottom edge and ends in the cell right before a W.
- **Only W forms a rectangle centered on the grid.** I checked every letter at every width from 3 to 48. Try rows of 38 in the figure: the same four W's still make a rectangle, but it sits off to one side.

The Half-Turn

Now turn the grid by 180° around its center cell: position p lands on position $96 - p$. **The two groups of the 2024 article exchange places, with four exceptions**:

- Segment 1 covers the first 20 letters of segment 6; the first nine letters of segment 3 cover segment 4; and segment 5 lands *exactly* on segment 2.
- The cribs overlap: BERLINCLOCK, read backwards, lands on the first eleven letters of EASTNORTHEAST.
- 44 letters of Group A land on 44 letters of Group B. Four letters have no partner in the other group: the ZZ at positions 47 and 48 (the second Z is the center cell itself, and the first lands on the central W) and the AR left over at the end.

[1] Section 4.1 of "A Vectorial Reading of Kryptos K4" notes the placement of both cribs against the outer W's (already observed in my 2024 article) and compares two reflection-symmetric anchor positions around the central W at position 49 (positions 25 and 73, with $25 + 73 = 98$). The grid here turns around the Z at position 48 ($p \mapsto 96 - p$). To my knowledge, there is no earlier public account of the four W corners in rows of 19 or the 44 pairs exchanged under that half-turn.

These symmetries are connected. The rectangle's corners pair around the center cell (positions 21 and 75, 37 and 59); together with the balanced groups, they produce the exchange described above.[2] The cribs overlap because one begins just after a corner and the other ends just before the opposite corner. The statistics below therefore count the rectangle and its half-turn symmetry as a single observation.

In the figure below, copies of the Group A letters flip top to bottom, then left to right, and come to rest on the Group B cells. Only AR remains uncovered.



There is an important caveat: **the grid holds 95 of K4's 97 letters, leaving the final AR outside it**. Its center is position 48; the center of the complete text is position 49. I have no explanation for AR. Treating it as filler or an end marker would be an assumption: a decryption must account for these two letters, as well as the unmatched ZZ near the center.[3]

The half-turn also gives something to test. Eleven pairs of opposite cells have known plaintext on both sides. One plaintext pair occurs twice: TL, at positions 25 and 71, then at positions 29 and 67. Their encryptions are VZ and RV, respectively. This rules out a fixed rule that encrypts the letters of opposite cells two by two.[4]

Just a Coincidence?

How often would this happen by chance? I updated K4nundrum, my open-source tool, to find symmetric rectangles in any ciphertext (`--rectangle`), count separator placements (`--rectangle-stats`), and simulate random ciphertexts (`--sim --rectangle`).

[2] If the five separator positions are `a < b < c < d < e`, equal alternating group lengths imply `a + c + e − b − d = 49`. The rectangle gives `a + e = b + d = 96`, which forces `c = 49`. The middle W's position therefore follows from the balance and the rectangle; it is not an additional independent coincidence.

[3] The Enigma manual cited in my 2024 article lists ZZ as a comma, beside X for a full stop. Here ZZ is ciphertext, not plaintext, so that convention does not apply directly. I note it because the pair sits just before the central W.

[4] One curiosity: from the first pair to the second, four cells away, both ciphertext letters move back by exactly four letters of the alphabet (V to R, Z to V). It is a single observation among the many comparisons I tried on these eleven pairs, and it may well be a coincidence.



K4nundrum, the open-source tool behind these measurements.

I first checked the 2024 pattern. Of all the ways to separate the six pieces into two groups, only the alternating split gives matching letter counts at every frequency. In a separate test, an anonymous researcher randomly split the same 92 letters into two groups of 46 a million times, and the counts matched in 0.45% of the trials (compared with under 2% in my 2024 simulation, which generated fresh random texts).[5]

Now for the rectangle. There are 64,446,024 ways of placing five separators among the 97 positions of K4. K4nundrum enumerates them all and counts those that show each pattern. Here, a *symmetric rectangle* means that, for some row width, the four outer separators sit on four cells that mirror one another left to right and top to bottom in the grid. (In the half-turn section, "opposite cells" meant the two cells that the half-turn exchanges.)

Pattern	Placements	About 1 in
Alternating groups of equal size	1,272,384	51
W-rectangle: the four outer separators on mirror cells, in rows of some width	26,734	2,411
W-rectangle, with the middle separator next to the center of the grid	1,754	36,742
W-rectangle, with its corners one cell in from the corners of the grid	775	83,156
Alternating groups of equal size and a W-rectangle	650	99,148

K4 shows all these patterns. Since I selected W in 2024 for its frequency pattern, a useful comparison is with separator placements that already produce two alternating groups of equal size. Of 1,272,384 such placements, **650 also form a rectangle: about 1 in 1,958** (1 in 1,582 when segments shorter than three letters are excluded, as K4nundrum does for K4-like groups). This figure is the product of two steps: in 35,720 of these placements the middle separator sits at position 49 (about 1 in 36), and among those, 650 admit a width that puts the four outer separators on a rectangle (about 1 in 55). These placement counts measure group sizes; the matching frequency shapes are tested in the simulation below.

Second, the simulation. With 26 letters to choose from, *some* letter forms *some* rectangle fairly often. In more than 119 million random 97-letter texts, generated by

[5] "A Vectorial Reading of Kryptos K4", an anonymous preprint dated 22 May 2026, section 4.1. Its test draws one million random splits of the 92 letters that are not W. It keeps K4's letters and changes only the half each one falls in, so it measures something different from the simulations below, which generate new random texts.

drawing each letter independently and uniformly from A to Z, K4nundrum found a letter forming:[6]

- a symmetric rectangle, at some width, in 4.6% of the texts (1.9% with 5 rows or more, 0.37% with a center mark);
- the same arrangement as K4, with corners one cell in and a center mark, in 0.010%, about 1 in 10,000;
- a rectangle, while also splitting the text into K4-like groups (the 2024 pattern), in only 49 texts, about 1 in 2.4 million;
- that same arrangement while also splitting the text into K4-like groups, as W does in K4: *not once in this run.*

> The fairest single figure is the conditional one: **among the 1,272,384 separator placements that already give two alternating groups of equal size, 650 also form a rectangle, about 1 in 1,958**. I noticed the shape before measuring its rarity, and I specified the inset corners and the center mark after seeing them. The smaller figures of the simulation (about 1 in 10,000 for this arrangement among random texts, about 1 in 2.4 million together with the 2024 frequency pattern) therefore describe the chosen random model, not the probability of a coincidence. The geometry is worth investigating; it does not prove that the pattern was intentional.

The simulation searches across letters and grid widths, but it does not account for every other pattern I might have noticed. And an encrypted message need not behave like uniformly random letters.

The W's Are Stops, Not Nulls?

In 2024, I concluded that the W's were not mapped to plaintext characters, leaving at most 92 characters of plaintext (97 minus the five W's). **This claim has to be revised.** In 2005, Ed Scheidt, who designed the Kryptos ciphers with Jim Sanborn, said that every letter in K4 will be used (notes from that meeting). In 2025, the New York Times reported that the plaintext found among Sanborn's papers has 97 characters, exactly as many as K4.

I now believe that each W stands for a character of the plaintext, **most likely a stop, such as X, or Q for a question mark**, as in K2 and K3. *The plaintext would then consist of six segments of 20, 15, 11, 9, 15 and 22 letters, separated by five stops.*

Did the cipher count the five W positions, or were the W's inserted afterwards? If they were inserted afterwards, they are foreign to the cipher's rhythm: removing them should make its periodic anomalies sharper, not weaker. Richard Bean lists three statistical anomalies involving periods of 7 or 21 in his 2021 paper "Cryptodiagnosis

[6] The K4nundrum README records the simulation and its definitions. "K4-like" requires alternating groups with matching frequency shapes and no segment shorter than three letters; two or more groups are allowed. The reported frequencies count texts in which at least one letter has the property, at any tested width from 3 to 48. The K4 arrangement means inset corners plus a center mark, across all the grid sizes searched; it does not require exactly five rows and nineteen columns. Zero occurrences in a finite run do not establish a zero probability.

of Kryptos K4"[7]. Let's remove the W's and count again. As a control, I also removed five letters picked at random, 300,000 times.[8]

The three anomalies, with and without the W's

Anomaly	With the W's (97 letters)	Without the W's (92 letters)	Random deletions giving an equally weak or weaker result
A letter repeated 7 positions later	9 pairs	7 pairs	77% of the time
Doubled letters starting in the same column, in rows of 7	5 of 6	2 of 7	44% of the time
Vertical pairs occurring twice, in rows of 21	11	0	2.6% of the time

The first anomaly weakens, the doubled letters disperse across the columns, and the repeated vertical pairs disappear. Removing the W of TWT, at positions 36 to 38, even creates a seventh doubled letter.

For the first two measures, deleting the W's is not especially unusual: five random deletions produce an equally weak or weaker result about three times out of four, and almost every other time, respectively. Losing all the repeated vertical pairs is less common, but still happens in 2.6% of the trials.

These comparisons show that deleting letters can disturb the patterns. **They do not establish when the W's entered the text.** I keep the idea that the cipher counted all 97 positions as a working hypothesis. One possibility is that the five cells were reserved in advance and filled with W's. Another is that the stops themselves were encrypted: a repeating key can contain the same letter in several places, so five identical W's do not rule that out.

A Prediction for K5

K5 may provide a test on another ciphertext. In his open letter of November 2025, Sanborn described another 97-character message, with a coding system similar to K4's and BERLINCLOCK in the same position. Paradigm, the firm that bought his archive, holds its ciphertext and says it plans to release it.[9]

A similar coding system would not necessarily preserve the same stop positions or geometry. Before K5's release, I would nevertheless look for two possible signs that the proposed structure carries over:

[7] Richard Bean, "Cryptodiagnosis of 'Kryptos K4'", *Proceedings of the 4th International Conference on Historical Cryptology (HistoCrypt 2021)*, Linköping University Electronic Press, 2021. His code is on GitHub. For the width-21 repeated pairs, he estimates about one occurrence in 6,750 random permutations of K4. That comparison preserves K4's letter counts and differs from generating new random texts or deleting five letters from K4.

[8] Each trial removes five distinct positions from K4 and recalculates the statistics on the remaining 92 letters. For the doublets, the comparison uses the largest number starting in any one column, not their proportion of all doublets. For the vertical pairs, it counts distinct pairs occurring at least twice. The percentages describe separate tests and should not be multiplied.

[9] See Sanborn's open letter of 12 November 2025 and Paradigm's June 2026 announcement. Sanborn specifies a similar coding system and BERLINCLOCK at the same position. The separator and rectangle tests above are my additional hypotheses.

- **The same positions:** a single letter, which need not be W, occupies positions 21, 37, 49, 59 and 75.
- **The same structure:** some letter forms a symmetric rectangle and splits the text into alternating groups with identical frequency shapes, as W does in K4.

K4nundrum can check both once K5 is available. If neither appears, that would weaken the idea that K4 and K5 share this structure. It would leave open whether the W's of K4 are stops, since a similar method could use different keys, stop positions or plaintext arrangements. If either appears, it would give a new lead to investigate, rather than a decryption by itself.

Conclusion

For now, I see the rectangle and the half-turn as constraints to explore. A convincing explanation must account for AR (the two letters left over) and ZZ, explain how the W positions are handled, and reproduce all 97 ciphertext letters from a coherent plaintext.

Series and Companion Articles

Kryptos K4 (glthr.com/k4)

#	Article	Address
1	Kryptos K4: The 'W' as a Separator	glthr.com/kryptos-k4-the-w-as-a-separator
2	**Kryptos K4: The W-Rectangle** (this article)	glthr.com/kryptos-k4-the-w-rectangle
3	Kryptos K4: Why the 'W' Matters	glthr.com/kryptos-k4-why-the-w-matters

How to cite this article

> Guillaume Lethuillier, "Kryptos K4: The W-Rectangle", glthr.com, November 11, 2026.
> https://glthr.com/kryptos-k4-the-w-rectangle

For BibTeX:

```
@misc{lethuillier2026kryptos,
  author = {Lethuillier, Guillaume},
  title  = {Kryptos K4: The W-Rectangle},
  year   = {2026},
  month  = nov,
  url    = {https://glthr.com/kryptos-k4-the-w-rectangle},
  note   = {glthr.com}
}
```

Your questions and leads are welcome: kryptos@glthr.com

This document was typeset with LaTeX from the article published at https://glthr.com/kryptos-k4-the-w-rectangle, which remains the reference version (as updated on November 28, 2026).