

# Coded Hashes of Arbitrary Images

(or: the last frontier of emoji encoding)

Steven R. Loomis  
srl@icu-project.org  
(individual contribution)

Keith Winstein  
keithw@cs.stanford.edu  
Stanford University

Jennifer 8. Lee  
jenny@jennifer8lee.com  
Emojination

2016-05-02  
<https://srl295.github.io>

## 1 Introduction

Emoji are pictographs (pictorial symbols) that are typically presented in a colorful cartoon form and used inline in text. [...] In Unicode 8.0, there is a total of 1,282 emoji, which are represented using 1,051 code points.<sup>1</sup>

Recently, there has been considerable interest in adding newly created pictorial symbols, not found in any existing character set, to the Unicode Standard as emoji.<sup>2</sup> Advocacy groups and others request these code points because Unicode plain text remains the dominant interoperable interchange format for messaging. In practice, before a new emoji can be used, a code point must be assigned and be recognized by the sender and receiver. The stated longer-term goal for Unicode is that implementations should support “embedded graphics, in addition to the emoji characters”.<sup>3</sup>

In this proposal, we describe a mechanism to uniquely identify arbitrary images within a plain-text Unicode character sequence. This will allow implementers to create their own emoji without needing to request and wait for the assignment of a code point. The basic idea is to encode a globally-unique *secure hash* of the emoji in a Unicode character sequence. Once the receiver knows the image’s hash, it may already have the corresponding image, or may have a choice of several mechanisms to retrieve it.

Our technique will gracefully degrade on legacy Unicode implementations, but our proposal is limited to allowing Unicode to *uniquely identify* an arbitrary image. We propose to leave to implementers the details of how to retrieve the actual image.

---

<sup>1</sup>Mark Davis and Peter Edberg, eds. *Unicode Technical Report 51: Unicode Emoji*. 2015. URL: <http://www.unicode.org/reports/tr51/>.

<sup>2</sup>*E.g.*, Taco Emoji campaign, Beard Emoji campaign, Dumpling Emoji campaign. There are currently 79 candidate emoji that have been assigned tentative code points.

<sup>3</sup>Davis and Edberg, *Unicode Technical Report 51: Unicode Emoji*, op. cit., Section 8, “Longer Term Solutions”.

## 1.1 Background

The demand for emoji characters has surged in the last few years as public has become enthralled by the multicolor glyphs that can appear in-line with text. Emoji have been the subject of an explosion of press pieces, numerous late-night television comedic segments, and even an upcoming movie from Sony Pictures Animation. Never has the encoding attracted such attention. As emoji use has proliferated, there have been increasing demands for better representation in the existing emoji set — ranging from hair variations (color and facial), to more female presence, to single-family households, to regional flags.<sup>4</sup>

However, the current emoji approval process takes upwards of 18 months, starting from when a proposal is first introduced, to when it moves out of the Unicode Emoji Subcommittee, to the Unicode Technical Committee, to ISO and back again for final approval. This current process, overseen by a relatively small organization whose stated main mission is focused on encoding existing languages, is widely understood as a bottleneck and also incredibly resource intensive on members' time.

To fill the demand for customizable glyphs, there has been an explosion of emoji-like images from a variety of vendors, ranging from personalized Bitmoji, to highly publicized “Kimoji”, to earnest refugee-themed emoji. These rogue “moji” are not technically emoji at all, but rather images or “stickers” that can be sent through messaging apps. In addition, vendors are using technical workarounds such as the Zero Width Joiner (ZWJ) and tagging to create emoji images that do not have to be assigned code points.

Meanwhile, the variants of single emoji created by different vendors leads to confusion in the intent of the sender and recipient.<sup>5</sup> The inconsistent imagery is at tension with Unicode's goal of having interoperable platform-neutral communication. A longer term architectural solution is needed.

## 2 Proposal

**Outline:** implementations will be able to create “implementation-defined emoji,” and allow their users to send them to receivers. The sender encodes a secure hash of the emoji (which serves as an unforgeable globally unique identifier) in a sequence of coded characters, which we call a Coded Hash of an Arbitrary Image (CHAI). Upon receiving the CHAI sequence, the receiving implementation may already know the corresponding image, or may have to request it either from the sender or from a third party.

---

<sup>4</sup>Mark Davis and Peter Edberg, eds. *Unicode Technical Report 52: Unicode Emoji Mechanisms*. 2016. URL: <http://www.unicode.org/reports/tr52/>.

<sup>5</sup>Hannah Miller et al. “‘Blissfully happy’ or ‘ready to fight’: Varying Interpretations of Emoji”. In: *ICWSM '16* (2016). URL: [http://grouplens.org/site-content/uploads/ICWSM16\\_Emoji-Final\\_Version.pdf](http://grouplens.org/site-content/uploads/ICWSM16_Emoji-Final_Version.pdf).

## 2.1 Representing the emoji

The first step to creating an implementation-defined emoji will be to represent the emoji in a canonical form. This format will need to be specified, but we do not do so here. As a strawman, we suggest a JSON document with three keys: `content-type` (any IANA-registered MIME Media Type), `image` (a Base64-encoded suggested rendering of the emoji, interpreted as the format named in the `content-type` field), and `name` (name of the emoji, which may be used in fallback rendering for visually-impaired users). We call this representation the “emoji description.”

## 2.2 Secure hash

The implementation will identify the emoji by taking the SHA-256 hash of the emoji description. Given a secure hash value, it is intended to be intractable to find a different input that produces the same hash value. This is known as “second-preimage resistance”; SHA-256 is believed to have strong second-preimage resistance.

## 2.3 Code points allocated to express the secure hash

We propose to allocate 65,536 ( $2^{16}$ ) code points to represent bits of the secure hash. Each code point will express 16 bits of the secure hash. We envision allocating Plane 13, from U+D0000 through U+DFFFFD. (The two code points at the end of the plane are noncharacters.) These 65,534 code points will represent the 16-bit values  $0000_{16}$  through  $FFFD_{16}$ . In addition, we envision allocating U+EFFFFC to represent the value  $FFFE_{16}$ , and U+EFFFFD to represent the value  $FFFF_{16}$ . We refer to these 65,536 code points as “CHAI characters.” Each code point will have general category Cf.

In UTF-8, UTF-16, and UTF-32, every code point outside the BMP requires 32 bits to encode. As a result, the efficiency of this scheme will be 50% in each Unicode encoding scheme (16 bits of the hash will consume 32 bits in memory, on disk, or on the wire).

## 2.4 Encoding an implementation-defined emoji

To code a CHAI, the sender first encodes a “fallback” base character that most closely approximates the implementation-defined emoji. This will allow rendering to degrade gracefully on receivers that do not support CHAIs.

Next, the sender encodes the secure hash by appending between between one and sixteen CHAI characters. This will represent a prefix of the SHA-256 hash of the emoji description. The sender chooses how many bits of the hash to include. (We expect 80 bits, or five CHAI characters, to be sufficient for typical use today. The risk of including too few bits is that the hash may no longer be globally unique, especially in the presence of an attacker who wishes to create a different image with the same hash prefix.)

Therefore, the actual encoding is:

`<basechar> + <CHAI char> + <CHAI char> + <CHAI char> ...`

where the *basechar* represents a fallback base character, and each CHAI character (from the range U+D0000 .. U+DFFFFD and U+EFFFC .. U+EFFFD) represents 16 bits of the secure hash of the emoji description.

## 2.5 Receiving and rendering a CHAI

On receipt, the receiver determines if it already knows of an emoji whose hash matches the prefix encoded in the CHAI characters. If so, it displays the emoji (either using the image provided in the emoji description, or an image known locally). Otherwise, the receiver displays the base character while it attempts to retrieve an emoji description whose hash matches the encoded hash prefix.

We do not specify how this retrieval will take place; we expect it to vary based on the messaging protocol. In some protocols, it might be easiest for the receiver to simply ask the sender to send the full emoji description. In others, the receiver might consult an online repository run by the implementing vendor. Alternately, the receiver might consult a community repository where anybody can contribute any emoji description.

Because the CHAI serves as an unforgeable globally unique identifier for the emoji description, the receiver is intended to have confidence that if it finds a matching emoji description from anywhere, it can display it per the sender’s intent.

Depending on the protocol, as the input arrives, the receiver may have some ambiguity about when the sequence of CHAI characters ends. A receiver may choose to wait until the next non-combining character (signaling the end of the combining character sequence), or a protocol-defined end-of-message signal, before retrieving the emoji description.

## 2.6 Privacy

As Unicode Technical Report #51 notes: “There are also privacy aspects to implementations of embedded graphics: if the graphic itself is not packaged with the text, but instead is just a reference to an image on a server, then that server could track usage.”<sup>6</sup>

We agree. A similar exposure could occur if the protocol provides for the receiver to ask the sender to send an unknown emoji description; this could unwittingly expose if anybody *else* has ever sent the same emoji to the receiver.

It is not the goal of this document to address all aspects of embedded graphics, and we intend to leave to domain experts and implementers the decision on how to handle these questions. For example, a client might behave similarly to current practice in retrieving external images in email. If the receiver already knows of the identified emoji, it can display it, but otherwise the receiver would display the fallback (base) character and ask the user to push a button before retrieving the emoji description.

---

<sup>6</sup>Davis and Edberg, *Unicode Technical Report 51: Unicode Emoji*, op. cit., Section 8, “Longer Term Solutions”.

## 3 Example

### 3.1 chai.json

Example Canonical Emoji Description

```
{"name": "CHAI", "content-type": "image/png", "image": "iVBORw0KGgoAAAANSUHEUgAAAEAAAABACAYAAACQaXHeAAACT01EQVR42u2ZMUvDQBIGS4YM4tChU/EHiFMRp450pTg4dBR/gENnp0IH8QdIR21TTdfqo0BQRRAEG8VN0nVy6iTFsX7epReJIU2rufQu8D7wUS5pL+/39vL17pLJAAAAAAAAAAAAAYxyyKZEa1bo1QqxqKbEgKrQK92AQoMGRhwGMSnSYECb1S2UuZAXuy01tjcSM65ZHXPpM80GnK6KzmS9wzoa65AXVZI7UWkjBvjzQ3YCTDgNqMf/sqBXWAhE6pyRfEsXFKDniR1XzBVwN4bPJPcizZMwbxzuKNnt961s9c1pdctoW9VsN9nnM2ofs3AFr75NztksvrW16am6R01ynVytPD6erpluUcJ1GHPe0Am0SRvwcMwzjhF/QvTAXwIVvQRed7q2u65QLpgLnyZguQnxxHiCPNFPwm0ZRn71WxR3hPIf1wPts09MQkXcnehFddhYT3038juB53k/v/E8GcFoyDgJmMcB06DaINaMbc4Ve064BwzkGFJdiwD9GhNBbjGPAMJAORRHVXL0J0ckXZTyhGiEGDE0qqjvUNDPAlmFA0dBjyfcKpM0T5TcBvOH/ySuASuiozvfMcN3a3hrqgamBpzImE1VfMmmZqWkVrmhoUnyhtC2Lms6mfOZUImacpiwE5SU/T7wFTYK4pesbZrXIBEF2kroJg2ArWAAdKkDqpbpcOMWPPGzc/Wk+LkTZVbde7mo2ID1G7WutvPig1Qu10/UFOHVL+wqWpigLJXdtM9CukKY5sBgAAAAAAAAAAAAAAAAAP7ANw53SLcSCEY2AAAAAE1FTkSuQmCC"}
```

### 3.2 chai.png

Sample PNG



Sample encoding

```
# Read chai.png, wrote chai.json
# SHA-256 hash:
37d8c5403d29ec7d6f59b02690414de77b0597735a953c2a00fba50b5f79bb5
# Going to use 5 CHAI chars or 80 bits.

# <U+2615> (base)
# <U+D37D8> IMAGE HASH 37D8
# <U+DC540> IMAGE HASH C540
# <U+D3D29> IMAGE HASH 3D29
# <U+DEC7D> IMAGE HASH EC7D
# <U+D6F59> IMAGE HASH 6F59
```

Sample decoding

```
U+2615 (non-CHAI)
U+D37D8 hash bits: 37D8
U+DC540 hash bits: C540
U+D3D29 hash bits: 3D29
U+DEC7D hash bits: EC7D
U+D6F59 hash bits: 6F59
# CHAI Hash: 37D8C5403D29EC7D6F59
```

The resulting hash matches the first 80 bits of  
37d8c5403d29ec7d6f59b02690414de77b0597735a953c2a00fba50b5f79bb5

## 4 Character Data

This proposal requests a total of 65,536 new characters to be encoded.

### 4.1 Character Properties

```
D0000;IMAGE HASH 0000;Cf;0;BN;;;;;N;;;;;
...
DFFFD;IMAGE HASH FFFD;Cf;0;BN;;;;;N;;;;;
EFFFC;IMAGE HASH FFFE;Cf;0;BN;;;;;N;;;;;
EFFFD;IMAGE HASH FFFF;Cf;0;BN;;;;;N;;;;;
```

### 4.2 Line Breaking

```
D0000..DFFFD;CM # Cf [65,534] IMAGE HASH
EFFFC..DFFFD;CM # Cf [2] IMAGE HASH
```

## References

- Davis, Mark and Peter Edberg, eds. *Unicode Technical Report 51: Unicode Emoji*. 2015. URL: <http://www.unicode.org/reports/tr51/>.
- eds. *Unicode Technical Report 52: Unicode Emoji Mechanisms*. 2016. URL: <http://www.unicode.org/reports/tr52/>.
- Miller, Hannah et al. “‘Blissfully happy’ or ‘ready to fight’: Varying Interpretations of Emoji”. In: *ICWSM '16* (2016). URL: [http://grouplens.org/site-content/uploads/ICWSM16\\_Emoji-Final\\_Version.pdf](http://grouplens.org/site-content/uploads/ICWSM16_Emoji-Final_Version.pdf).

## Colophon

Repo URL: <https://github.com/srl295/srl-unicode-proposals>

Typeset by L<sup>A</sup>T<sub>E</sub>X. Made with 100% recycled bits. All opinions belong to the authors and do not reflect the opinions of their associated employers. Thank you to the Computer History Museum for hosting the “Emoji (and more)” event, where this proposal was originally penned.