ANRW2024
Applied Networking Research Workshop

HTTP/3’s Extensible Prioritization Scheme in the Wild

Joris Herbots, Robin Marx, Maarten Wijnants, Peter Quax, Wim Lamotte
Expertise Center for Digital Media – Hasselt University

Tuesday, July 23, 2024
IETF 120 – Applied Networking Research Workshop 2024

maxR
Mixed Augmented and Extended Reality Media Pipeline

UHASSELT
KNOWLEDGE IN ACTION

The HTTP Trilogy

HTTP/1.1 (TCP) - 1997

1 resource/connection

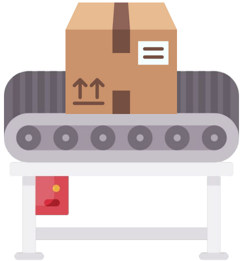

HTTP/2 (TCP) - 2015

≥1 resources/connection

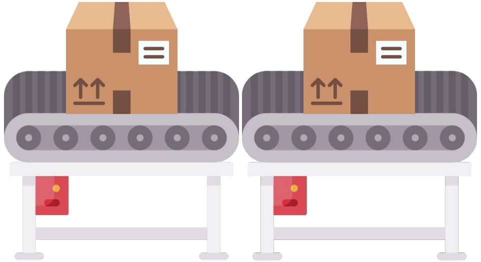

HTTP/3 (QUIC) - 2022

≥1 resources/connection

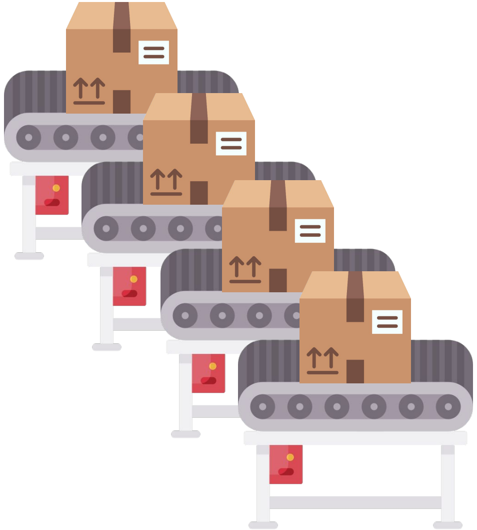

Images: Freepik

How do we indicate resource relationships?

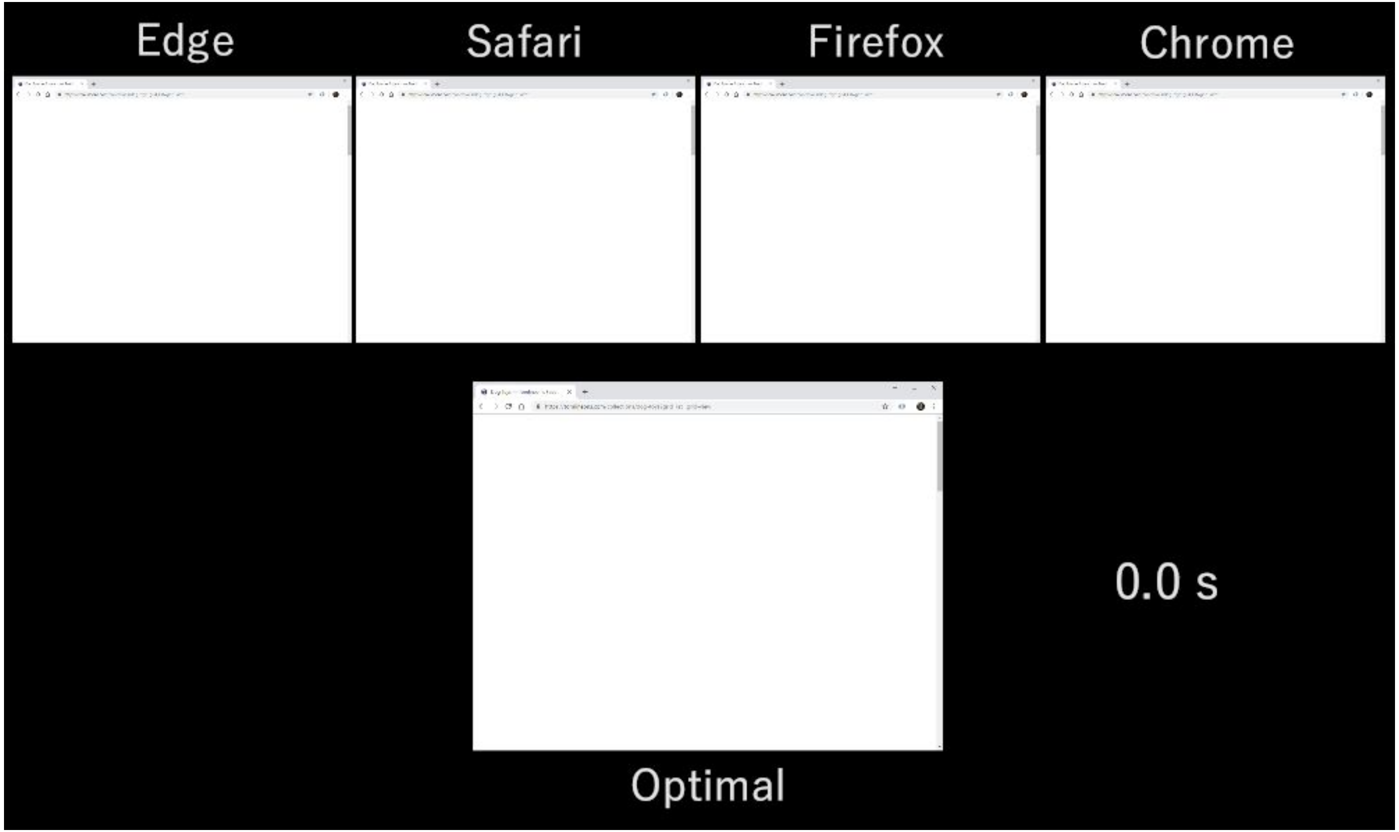

Priorities comparison image: https://blog.cloudflare.com/better-http-2-prioritization-for-a-faster-web

HTTP Priorities!

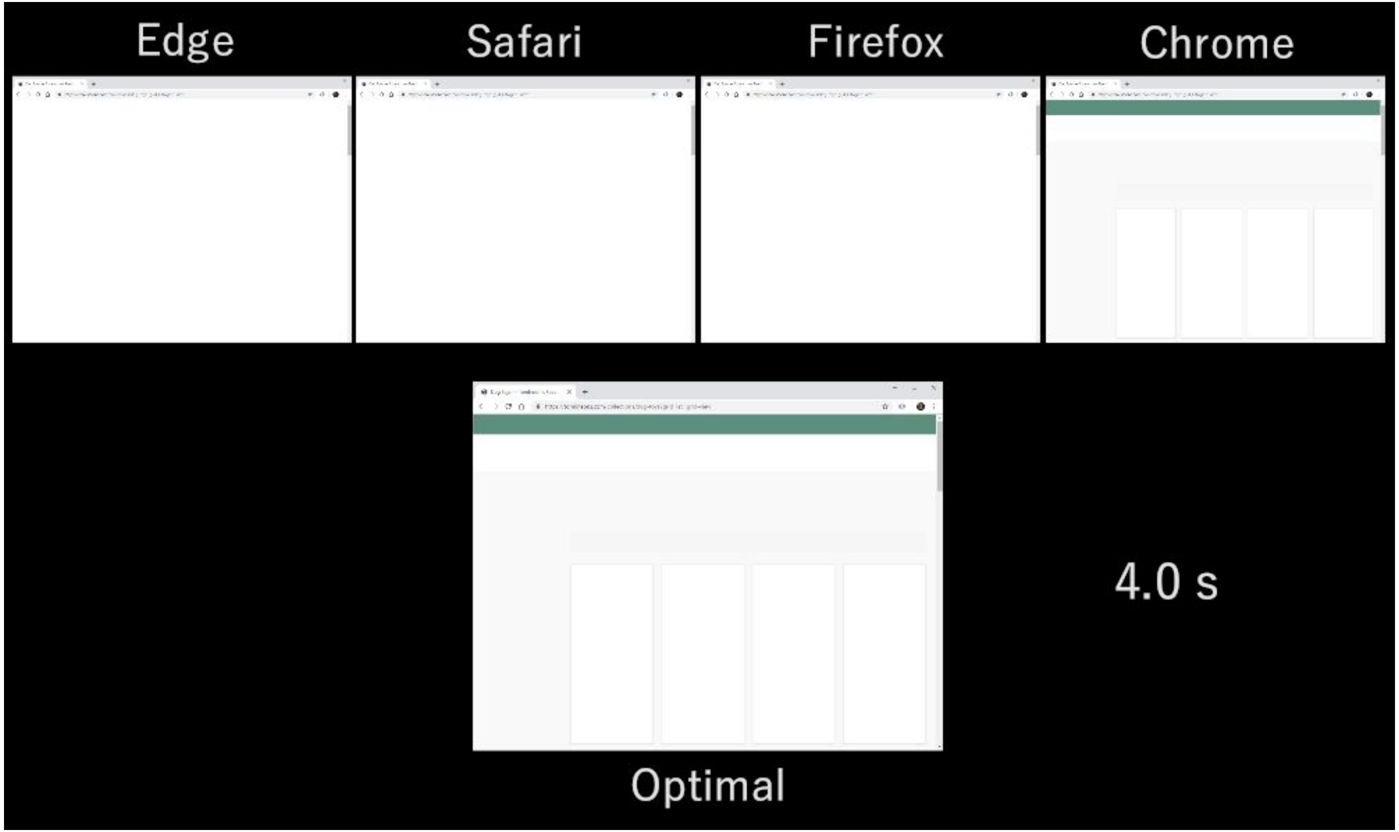

Priorities comparison image: https://blog.cloudflare.com/better-http-2-prioritization-for-a-faster-web

HTTP Priorities!

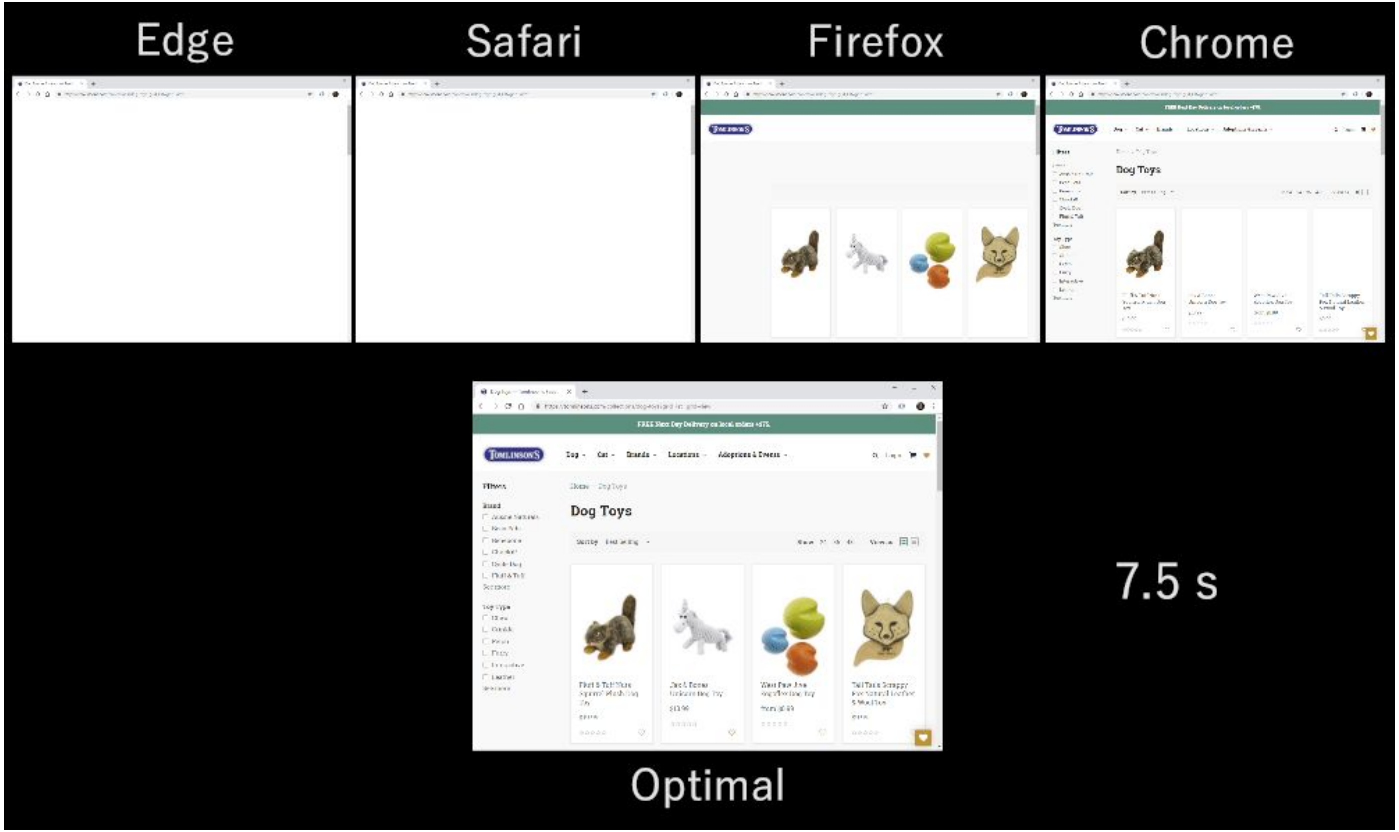

Priorities comparison image: https://blog.cloudflare.com/better-http-2-prioritization-for-a-faster-web

UHASSELT EDM

HTTP Priorities!

HTTP/2 Priorities

“It suffered from limited deployment and interoperability”

❗ Deprecated

Extensible Prioritization Scheme

Simple scheme of 2 parameters, extensible, can be backported to HTTP/2

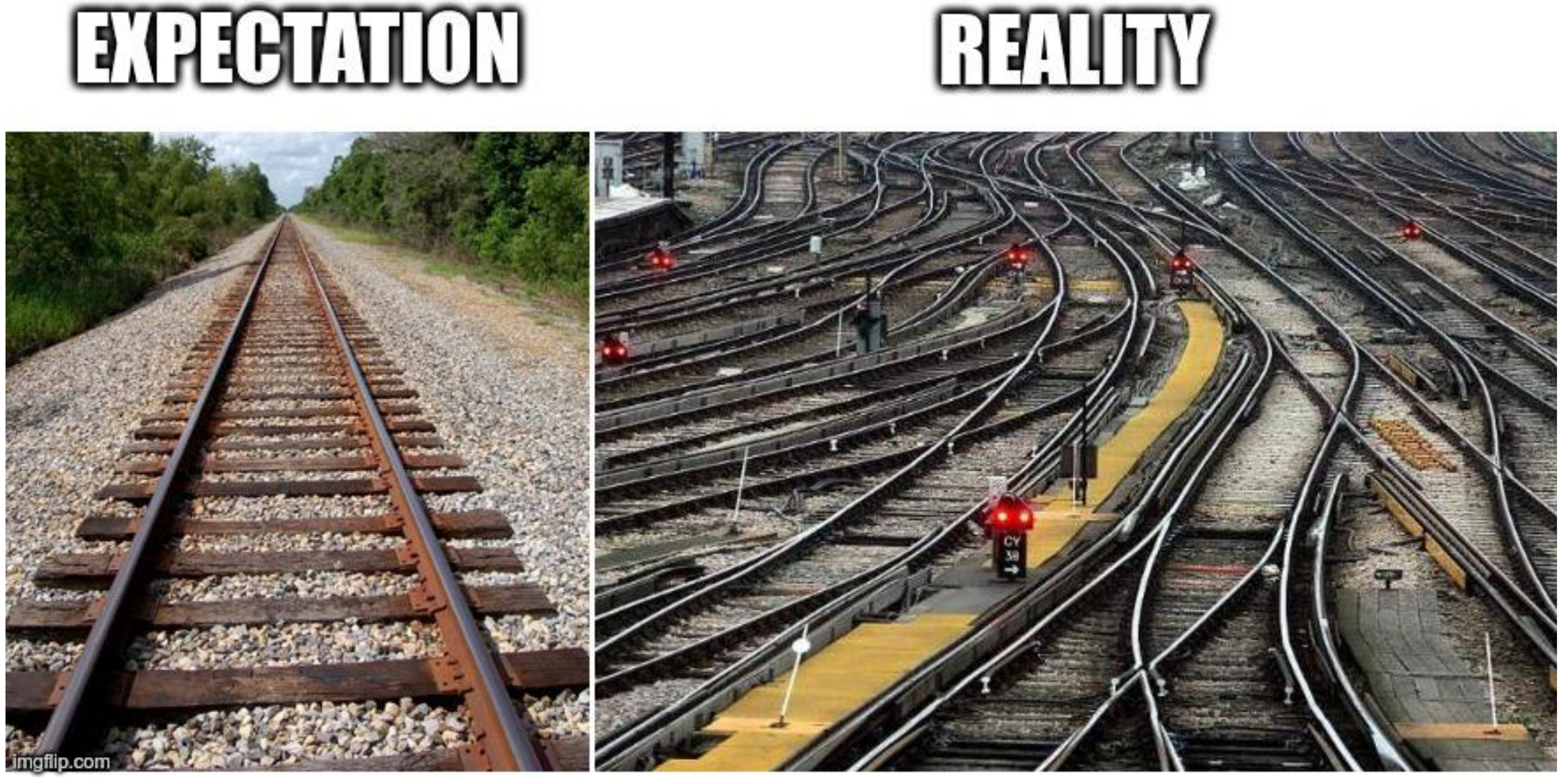
EXPECTATION
REALITY
imgflip.com
CY
38

Extensible Prioritization Scheme (EPS) in the wild

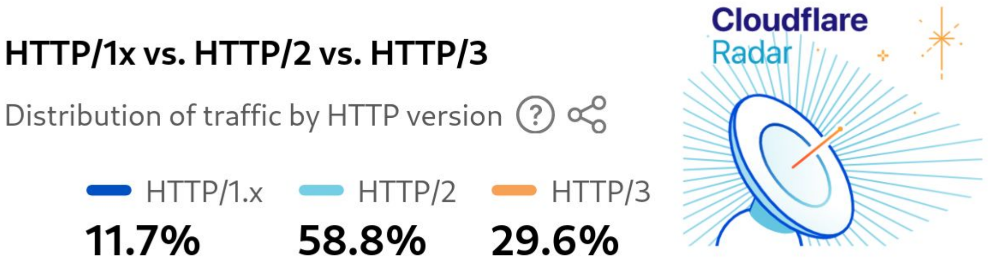

HTTP/1x vs. HTTP/2 vs. HTTP/3

Distribution of traffic by HTTP version

▬ HTTP/1.x	▬ HTTP/2	▬ HTTP/3
11.7%	**58.8%**	**29.6%**

First real world measurement study of EPS across prevalent HTTP/3 setups

- HTTP/3 and EPS are 2 years old
- ~30% web traffic is HTTP/3

Impact on web performance was a non-goal of this study

Browsers - Experimental Setup

Multiple HTML pages with a variety of resource types

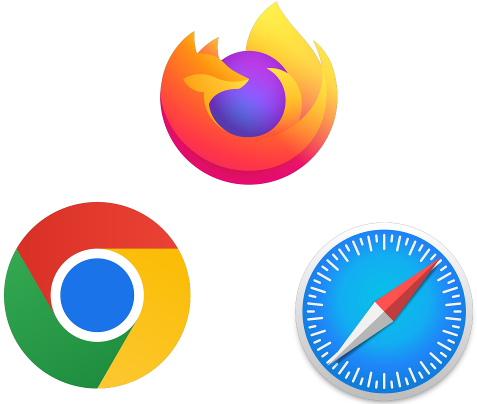

Custom aioquic server that recognizes EPS signals for logging

Experiments ran over a period of 1.5 years

- December 2022
- August 2023
- March 2024

EPS Parameters

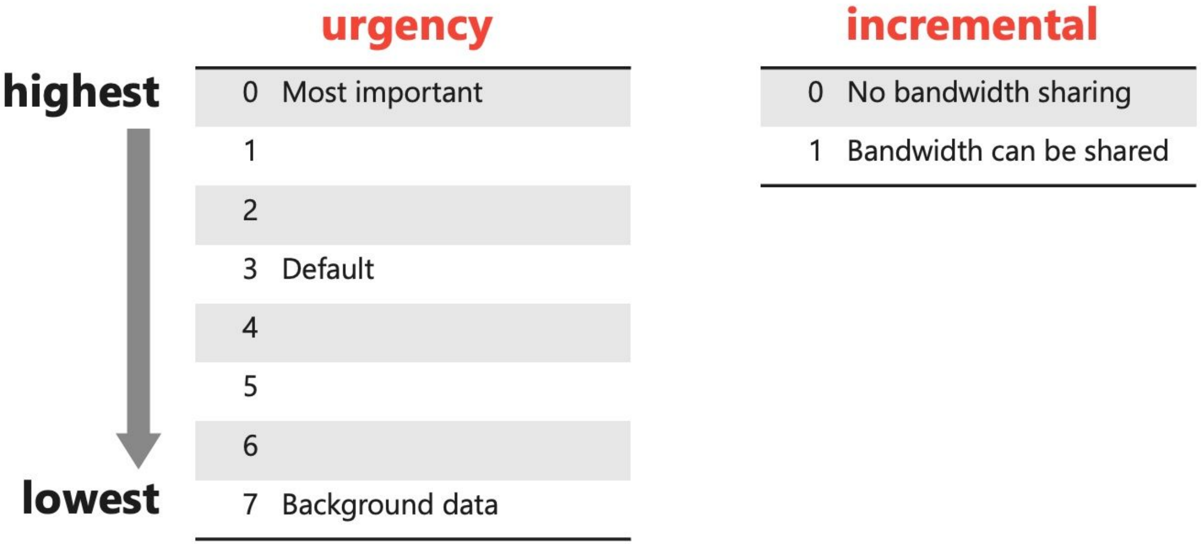

urgency

highest ↓ lowest

0	Most important
1	
2	
3	Default
4	
5	
6	
7	Background data

incremental

0	No bandwidth sharing
1	Bandwidth can be shared

HTTP header and/or binary frame

Browsers - Parameter Usage & Signalling

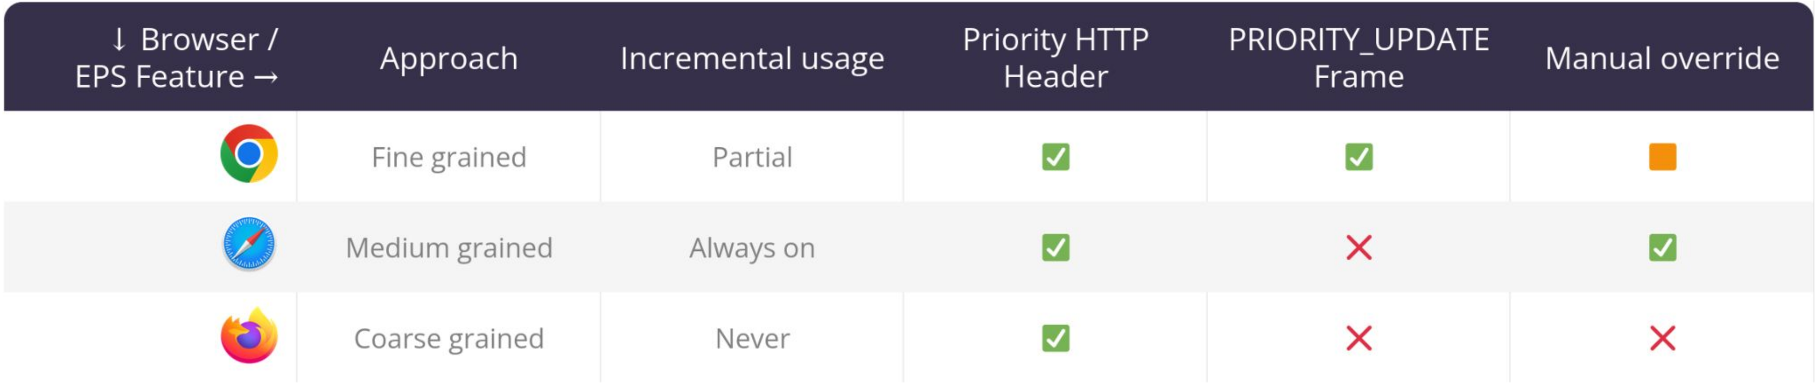

↓ Browser / EPS Feature →	Approach	Incremental usage	Priority HTTP Header	PRIORITY_UPDATE Frame	Manual override
	Fine grained	Partial	✓	✓	🟧
	Medium grained	Always on	✓	✗	✓
	Coarse grained	Never	✓	✗	✗

❗ A lot of heterogeneity

Browsers - Inconsistent Heuristics

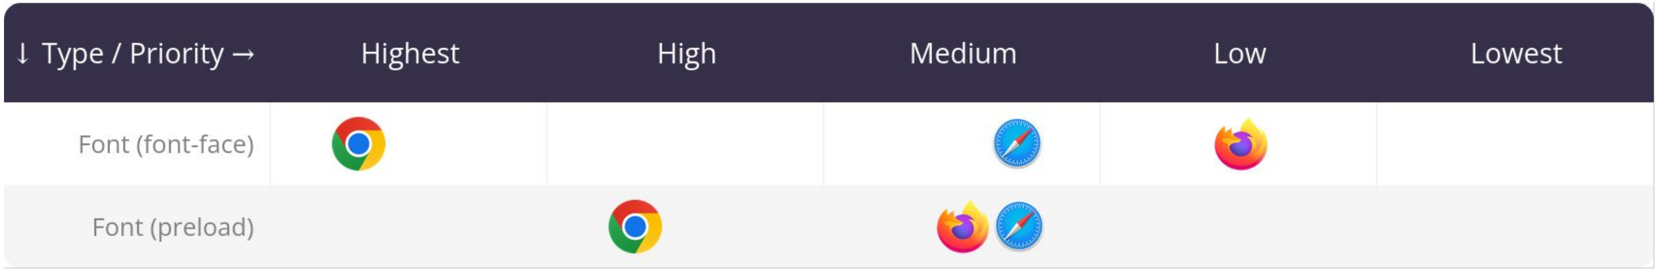

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
Font (font-face)	Chrome		Safari	Firefox	
Font (preload)		Chrome	Firefox Safari		

❗ Low priority custom font == Later arrival → Cumulative Layout Shifts

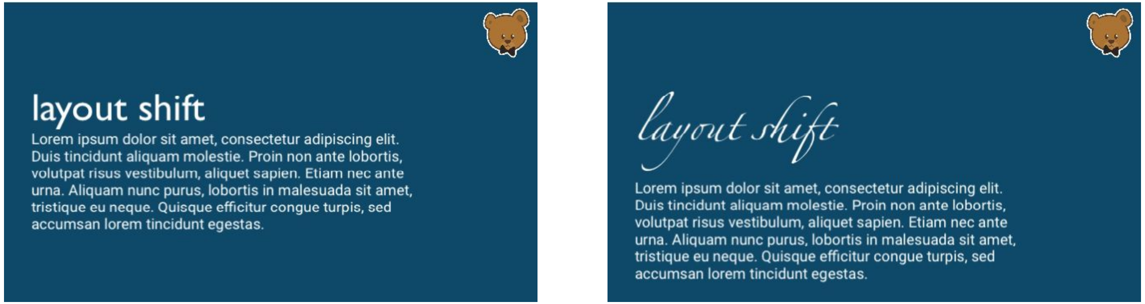

layout shift

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Duis tincidunt aliquam molestie. Proin non ante lobortis, volutpat risus vestibulum, aliquet sapien. Etiam nec ante urna. Aliquam nunc purus, lobortis in malesuada sit amet, tristique eu neque. Quisque efficitur congue turpis, sed accumsan lorem tincidunt egestas.

layout shift

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Duis tincidunt aliquam molestie. Proin non ante lobortis, volutpat risus vestibulum, aliquet sapien. Etiam nec ante urna. Aliquam nunc purus, lobortis in malesuada sit amet, tristique eu neque. Quisque efficitur congue turpis, sed accumsan lorem tincidunt egestas.

Layout shift image: https://www.debugbear.com/blog/web-font-layout-shift

Browsers - Inconsistent Heuristics

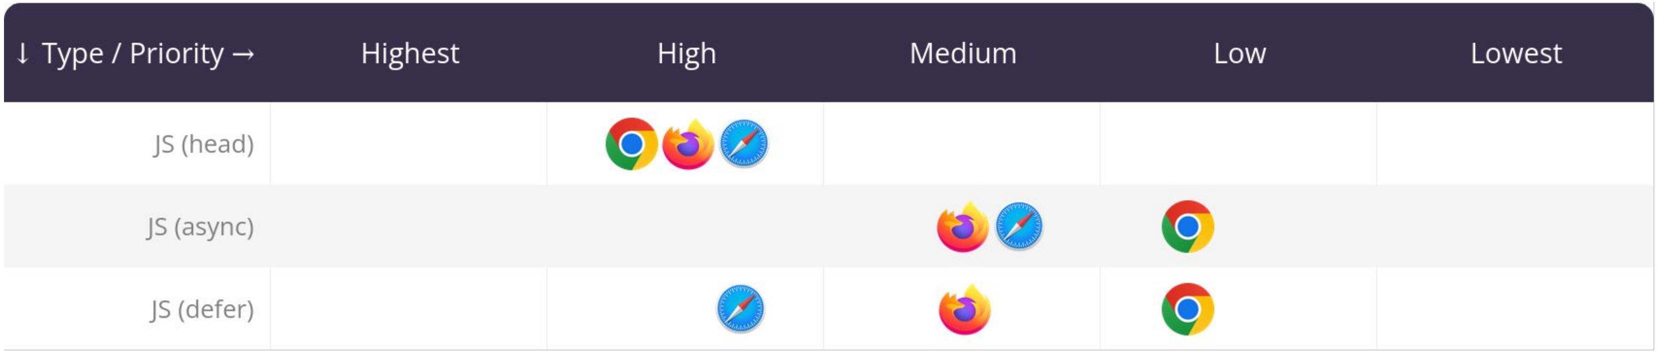

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
JS (head)		● ● ●			
JS (async)			● ●	●	
JS (defer)		●	●	●	

Browsers - Fetchpriority

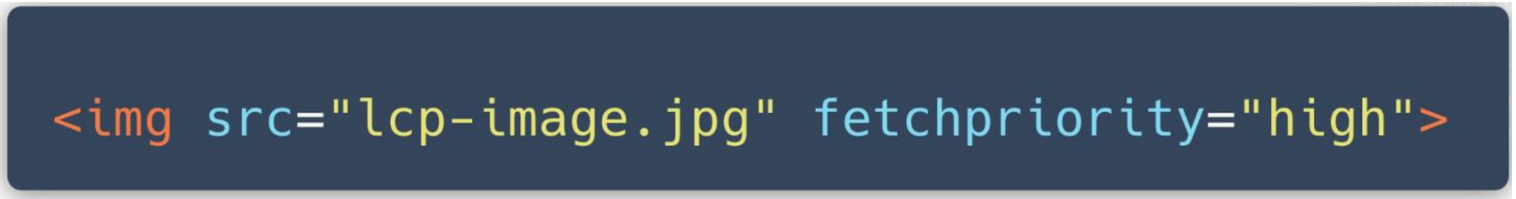

```html
<img src="lcp-image.jpg" fetchpriority="high">
```

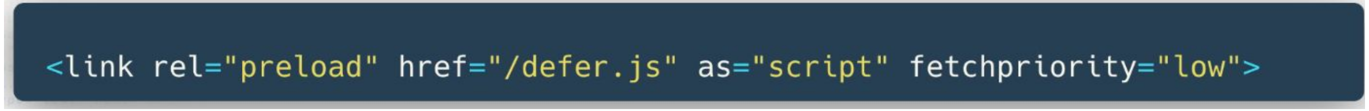

```html
<link rel="preload" href="/defer.js" as="script" fetchpriority="low">
```

Browsers - Fetchpriority

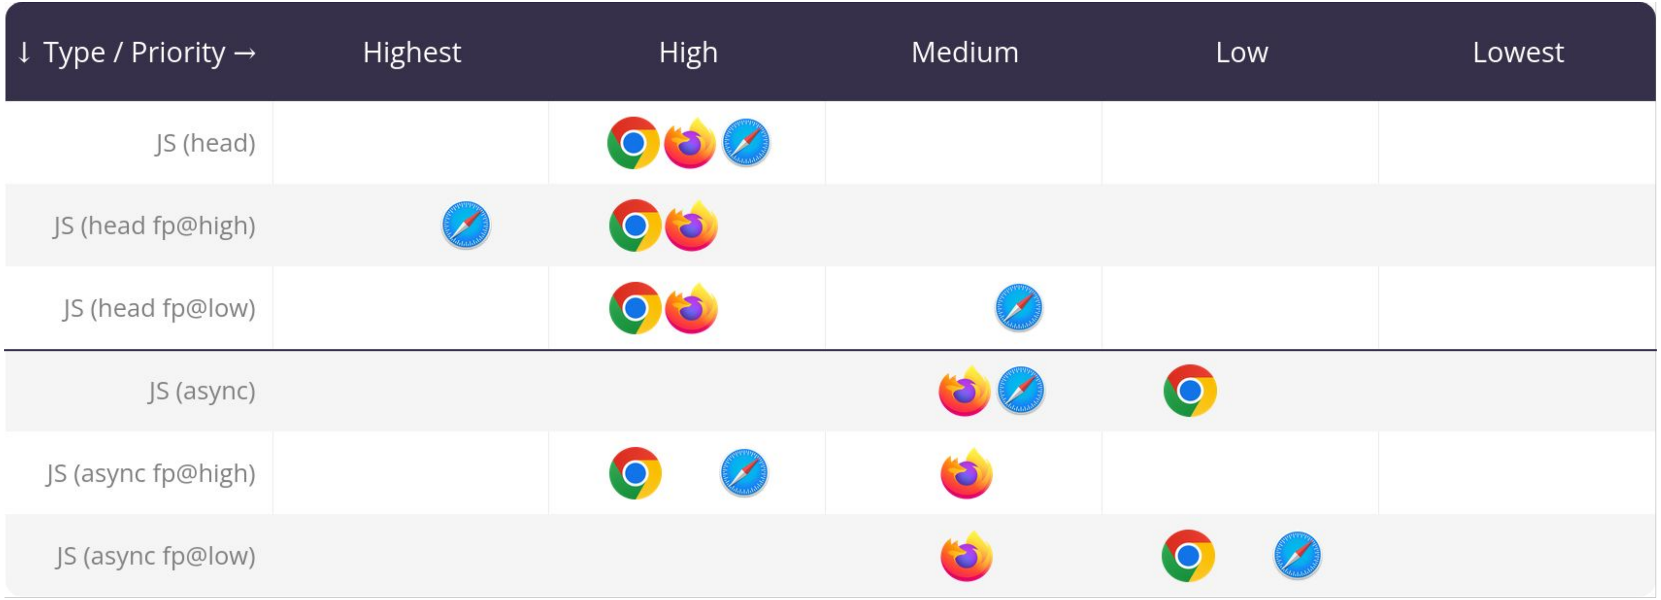

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
JS (head)					
JS (head fp@high)					
JS (head fp@low)					
JS (async)					
JS (async fp@high)					
JS (async fp@low)					

UHASSELT EDM

FETCHPRIORITY AND FETCH ALLOW FOR LIMITED STEERING
AND SERVERS DO EVERYTHING AS EXPECTED
SERVERS DO EVERYTHING AS EXPECTED, RIGHT?
imgflip.com

Servers - Experimental Setup

Modified aioquic client with EPS support

12 popular server stacks

Akamai
fastly
QUIC.cloud
amazon cloudfront
CLOUDFLARE
JSDELIVR
caddy
shopify
NGINX
Google Cloud

6 repetitions over 7 weeks

10.000 qlogs, manually analyzed with the qvis tools

Multiple resources in a wide variety of ways

Servers - Experimental Setup

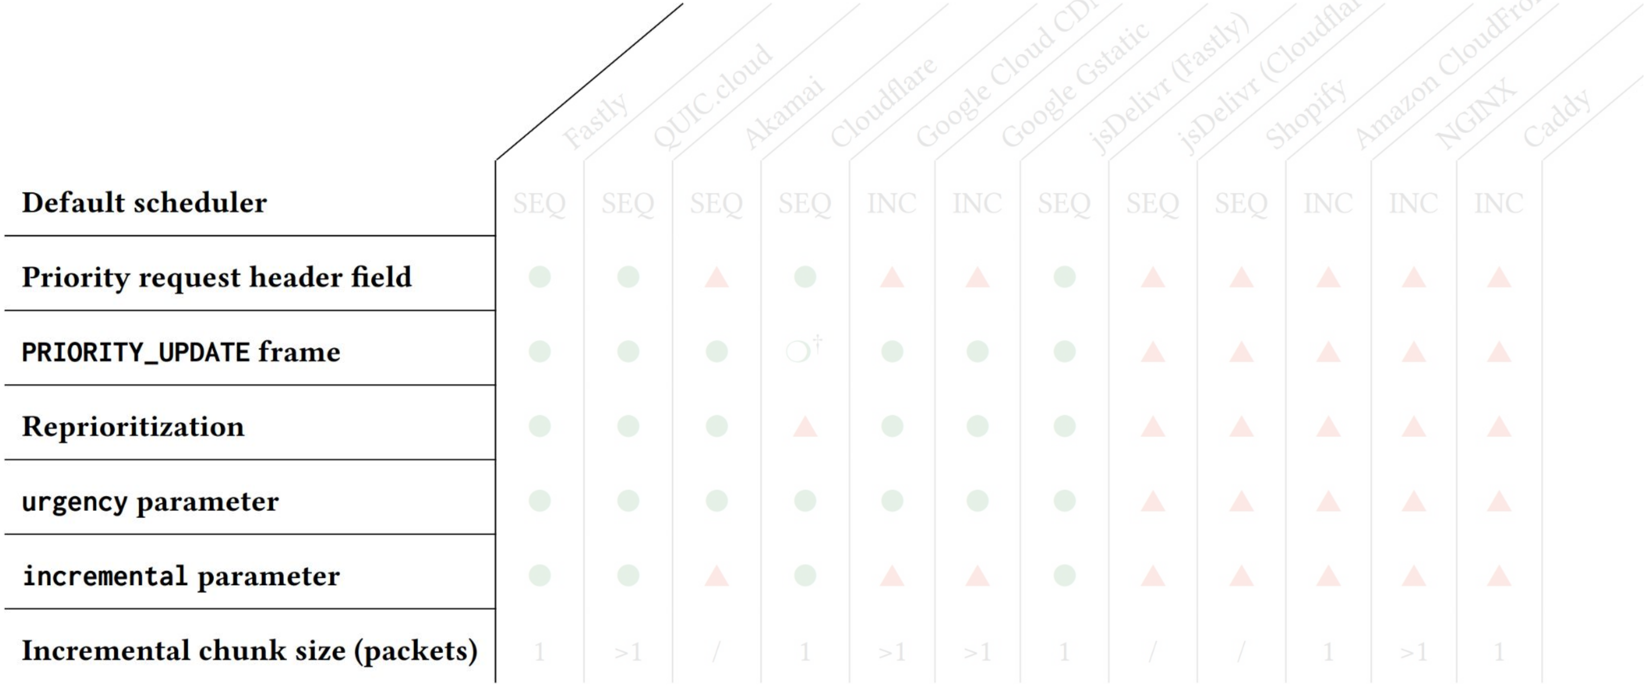

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
`PRIORITY_UPDATE` frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
`urgency` parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
`incremental` parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Servers - No Support

EPS support for NGINX and Caddy is on the roadmap

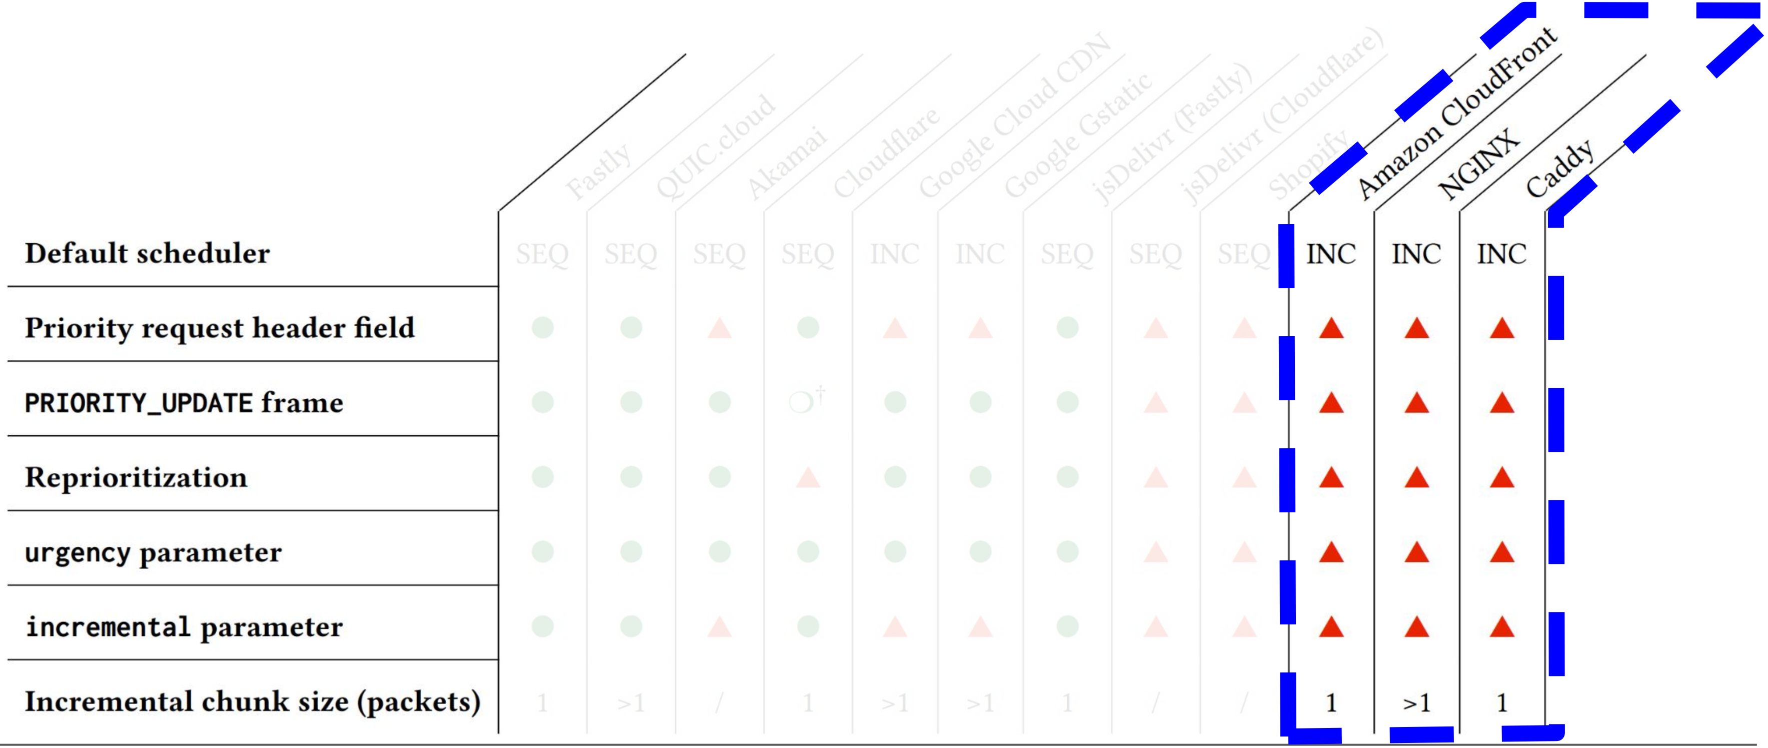

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Servers - Full Support

Minor difference in incremental chunk sizes

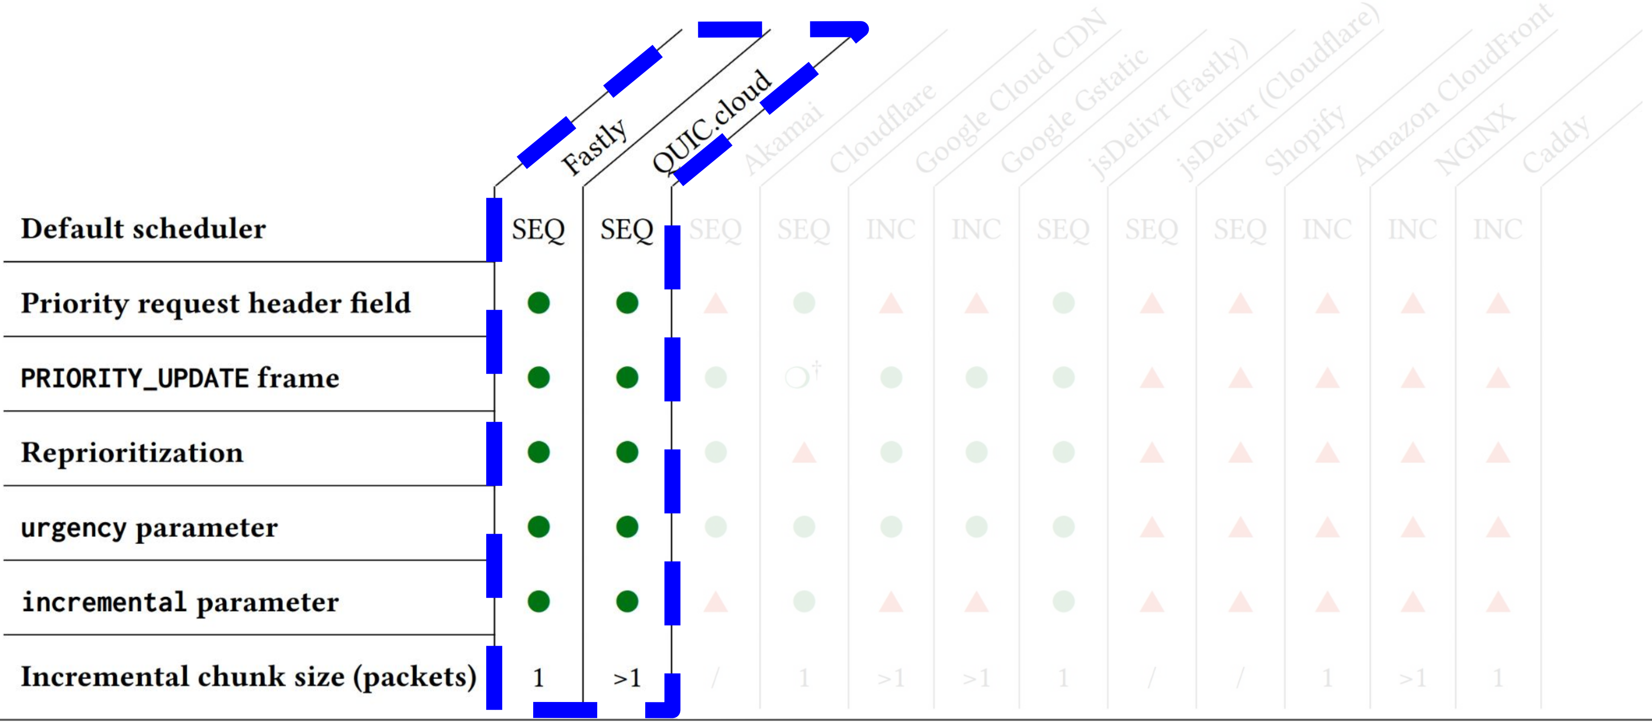

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Servers - Partial Support

Akamai and Google lack `Priority` header support

❗ Firefox and Safari request signals will be ignored

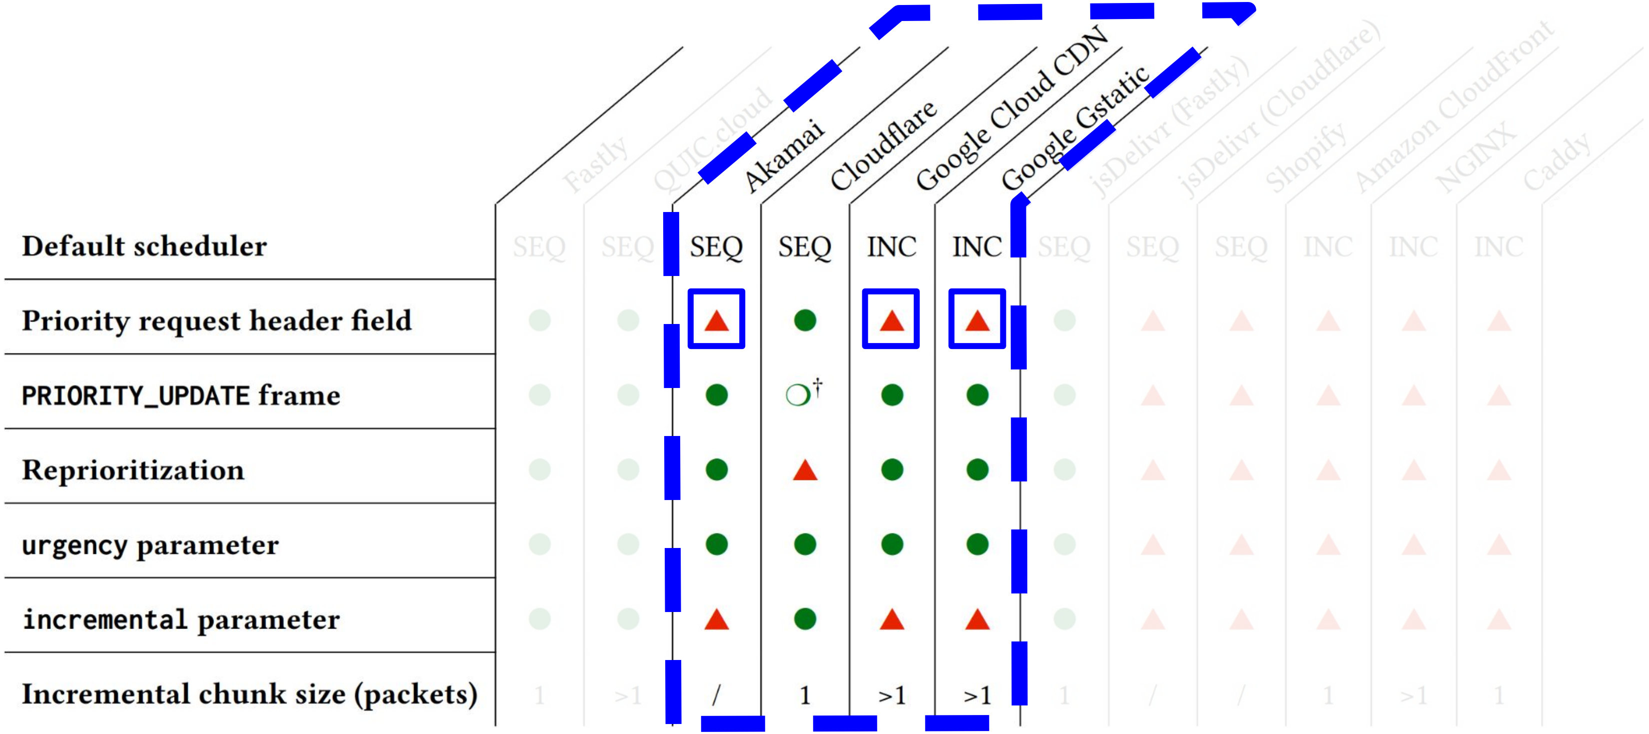

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Servers - Partial Support

Akamai and Google have no support for the incremental flag

❗ Google has no sequential support

❗ Chrome requests sequential loads

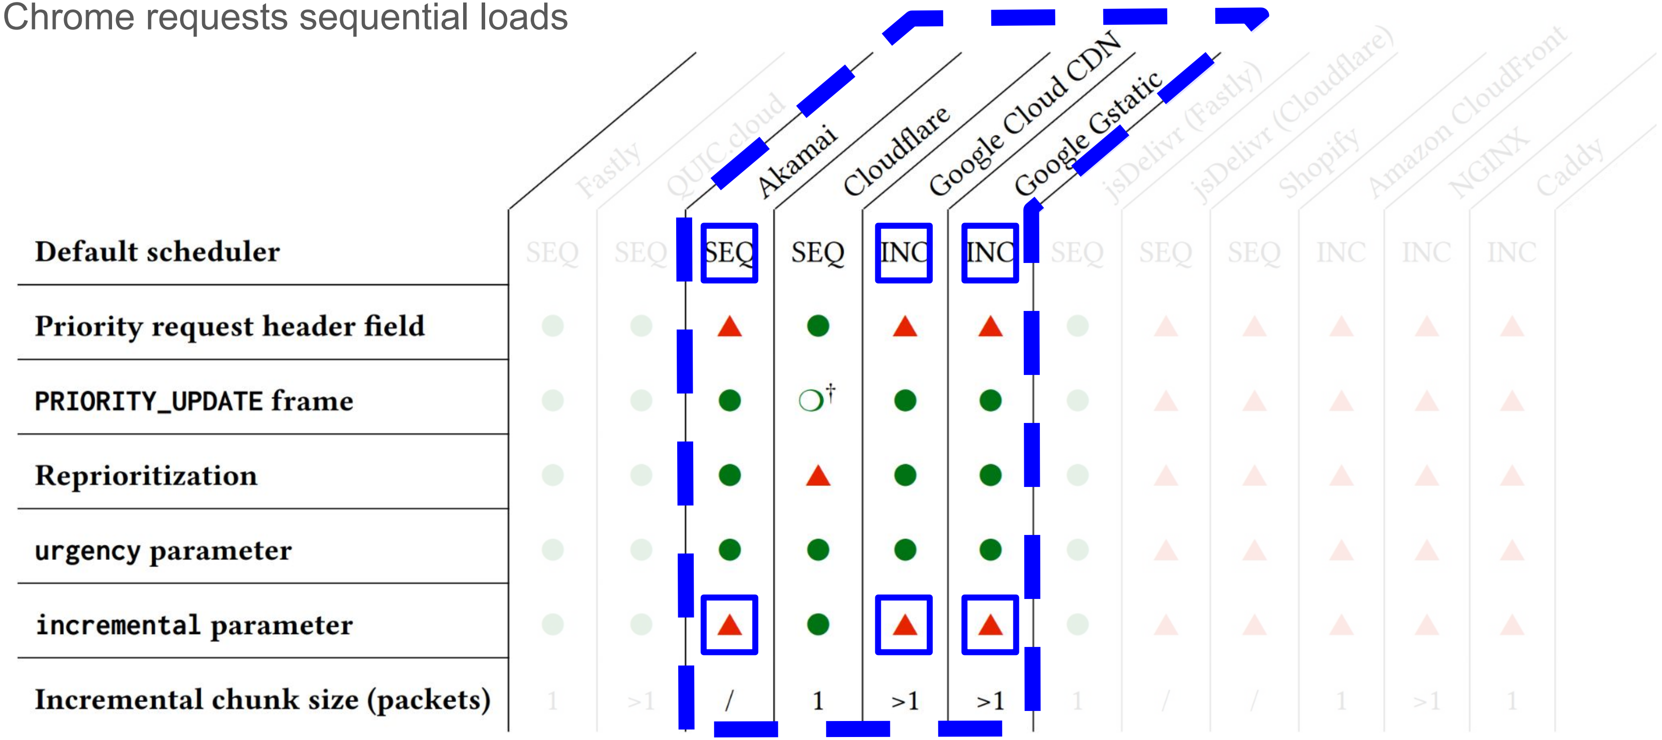

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Servers - Indirect Support

- ❗ Cloudflare has multiple stacks running in the same CDN?
- ❗ jsDelivr’s Fastly backend does support EPS → Inconsistent behavior

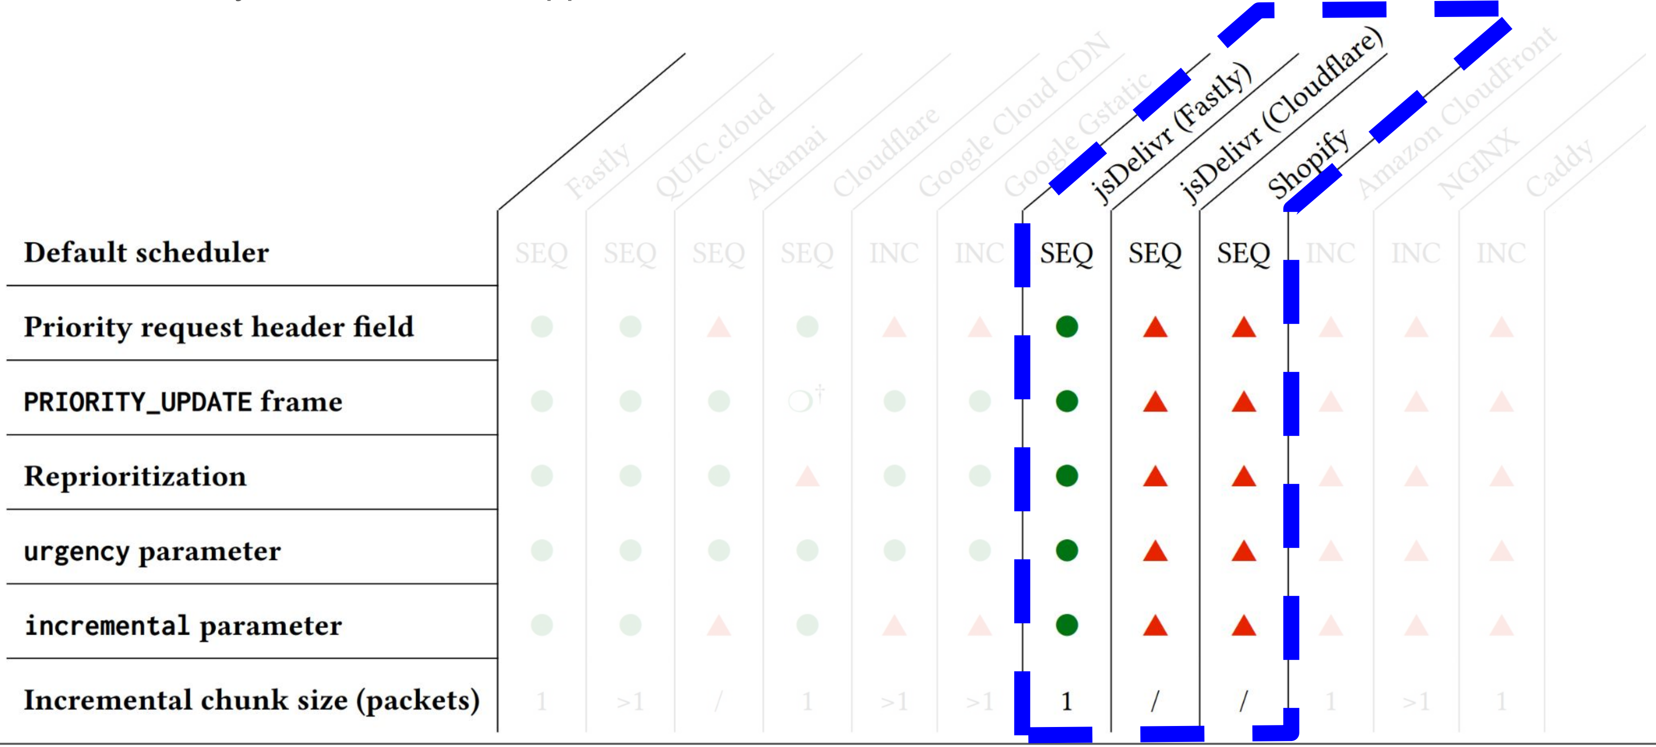

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

Considerable heterogeneity

Behavior unpredictable between browsers and servers

❗ `fetchpriority` ineffective

❗ Optimize for one browser, worsen experience in others

⇒ <u>Worst case!</u>

Priorities impact Web performance

❗ Websites/web developers care <u>a lot</u> about this

Recommendations

1. Full support for EPS by major deployments
    - Inconsistent without basic features!
    - Extensibility at risk
        - Space for future extensions
        - No basic support == no extensibility

2. Better manual control through developer APIs
    - Work around browser heuristics
        - E.g., change incremental through `fetchpriority`
    - Enable complex Web applications (e.g., video streaming)

Recommendations

3. Further research into loading heuristics and prioritization strategies
    - Browser engines use different prioritization strategies
    - Combat heterogeneity ⇒ Improve end-user experiences

4. Major deployments should offer realistic test resources
    - Ensure consistent testing conditions
    - Enable thorough validation of QUIC+HTTP/3 and its future features

Feel free to peruse our browser findings

You can access a detailed table online via the QR code (or the paper).

Extra Slides

EPS - Browsers - Main resource

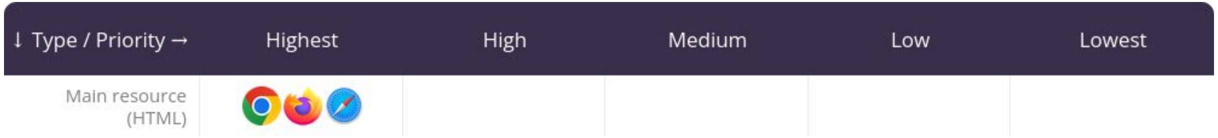

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
Main resource (HTML)	● ● ●				

EPS - Browsers - Fonts

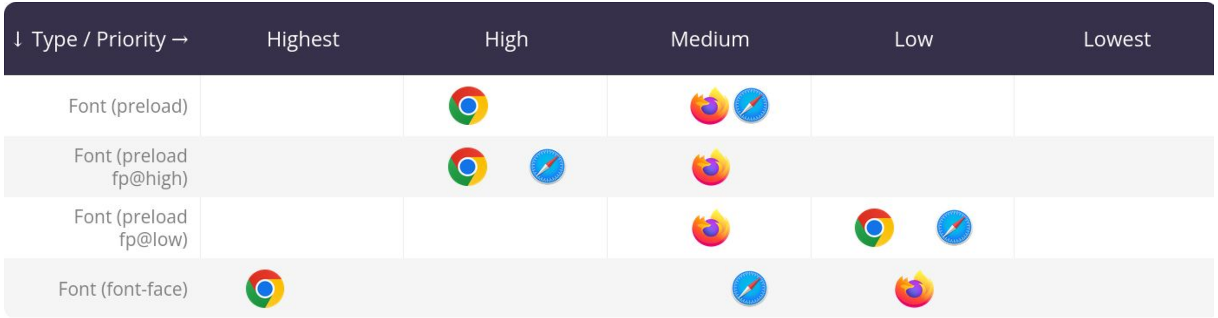

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
Font (preload)		Chrome	Firefox Safari		
Font (preload fp@high)		Chrome Safari	Firefox		
Font (preload fp@low)			Firefox	Chrome Safari	
Font (font-face)	Chrome		Safari	Firefox	

EPS - Browsers - JavaScript

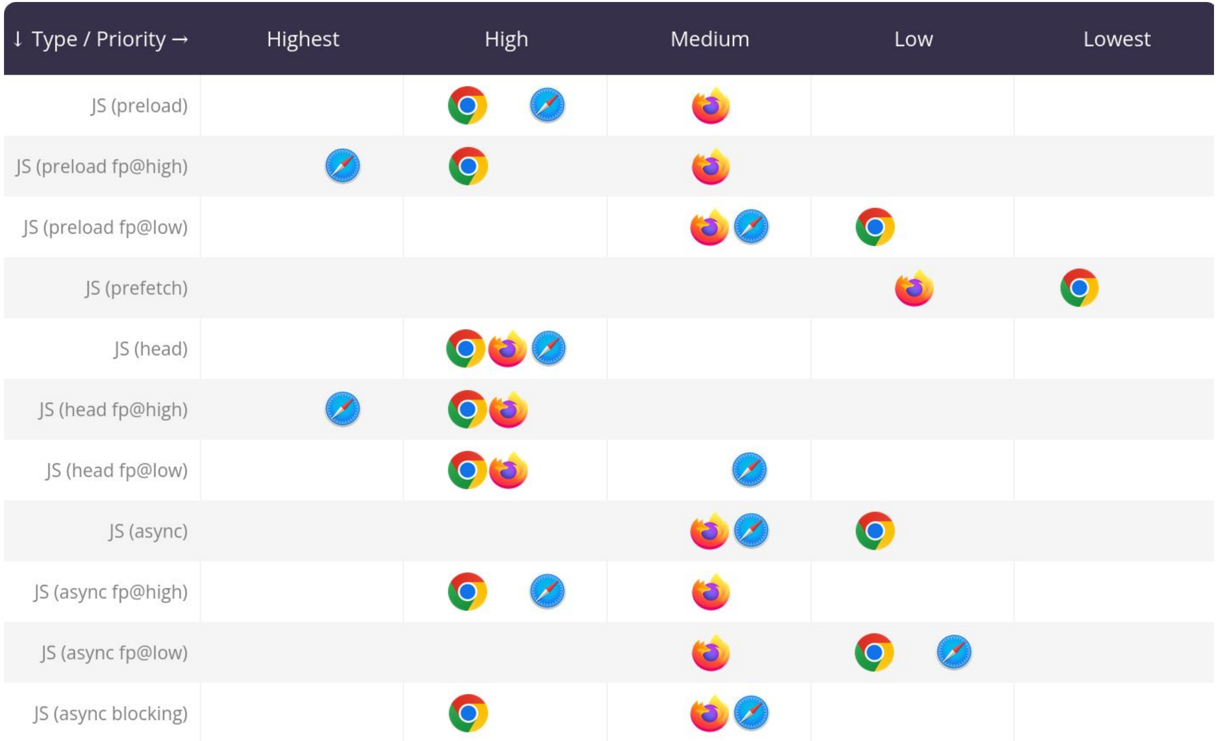

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
JS (preload)		◉ 🧭	🦊		
JS (preload fp@high)	🧭	◉	🦊		
JS (preload fp@low)			🦊 🧭	◉	
JS (prefetch)				🦊	◉
JS (head)		◉ 🦊 🧭			
JS (head fp@high)	🧭	◉ 🦊			
JS (head fp@low)		◉ 🦊	🧭		
JS (async)			🦊 🧭	◉	
JS (async fp@high)		◉ 🧭	🦊		
JS (async fp@low)			🦊	◉ 🧭	
JS (async blocking)		◉	🦊 🧭		

EPS - Browsers - JavaScript

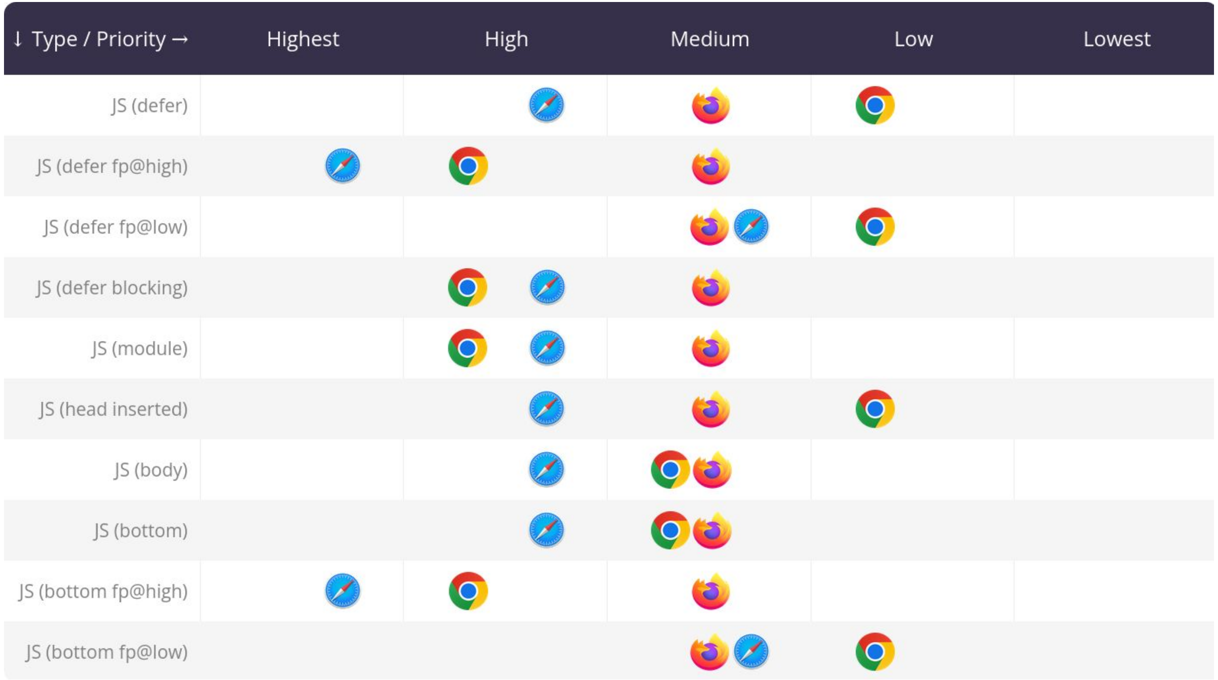

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
JS (defer)		Safari	Firefox	Chrome	
JS (defer fp@high)	Safari	Chrome	Firefox		
JS (defer fp@low)			Firefox Safari	Chrome	
JS (defer blocking)		Chrome Safari	Firefox		
JS (module)		Chrome Safari	Firefox		
JS (head inserted)		Safari	Firefox	Chrome	
JS (body)		Safari	Chrome Firefox		
JS (bottom)		Safari	Chrome Firefox		
JS (bottom fp@high)	Safari	Chrome	Firefox		
JS (bottom fp@low)			Firefox Safari	Chrome	

EPS - Browsers - CSS

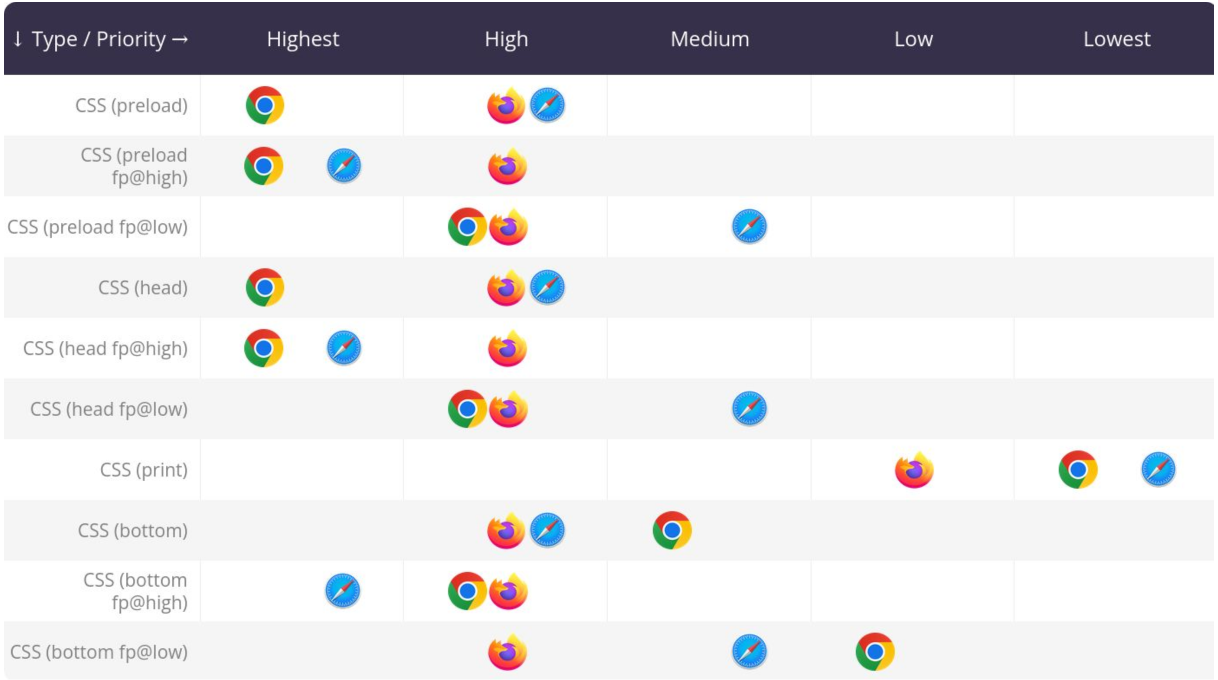

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
CSS (preload)					
CSS (preload fp@high)					
CSS (preload fp@low)					
CSS (head)					
CSS (head fp@high)					
CSS (head fp@low)					
CSS (print)					
CSS (bottom)					
CSS (bottom fp@high)					
CSS (bottom fp@low)					

EPS - Browsers - Images

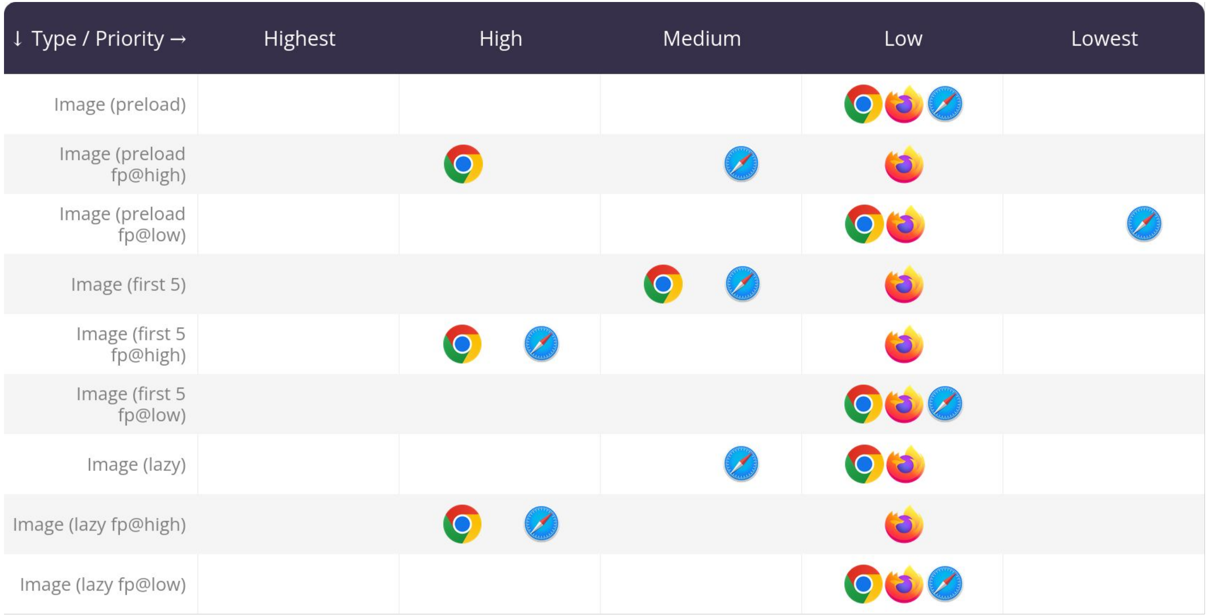

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
Image (preload)				Chrome Firefox Safari	
Image (preload fp@high)		Chrome	Safari	Firefox	
Image (preload fp@low)				Chrome Firefox	Safari
Image (first 5)			Chrome Safari	Firefox	
Image (first 5 fp@high)		Chrome Safari		Firefox	
Image (first 5 fp@low)				Chrome Firefox Safari	
Image (lazy)			Safari	Chrome Firefox	
Image (lazy fp@high)		Chrome Safari		Firefox	
Image (lazy fp@low)				Chrome Firefox Safari	

EPS - Browsers - Fetch

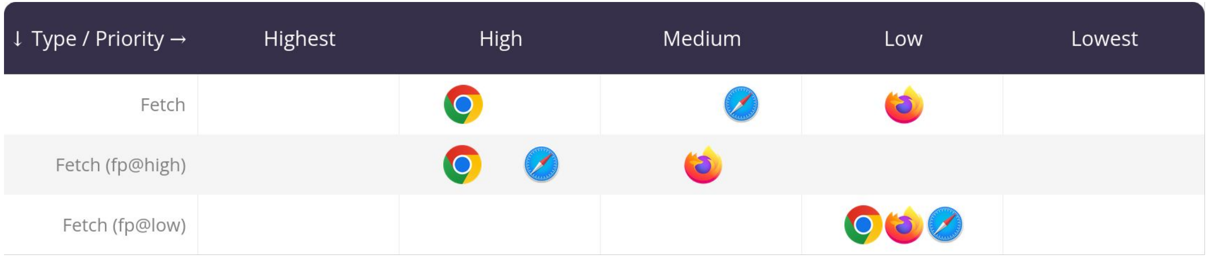

↓ Type / Priority →	Highest	High	Medium	Low	Lowest
Fetch		Chrome	Safari	Firefox	
Fetch (fp@high)		Chrome Safari	Firefox		
Fetch (fp@low)				Chrome Firefox Safari	

EPS - Servers - Partial Support

Cloudflare accepts `PRIORITY_UPDATE` frames BUT

- ❗ Ignores them after the request has come through
- ❗ No support for reprioritization

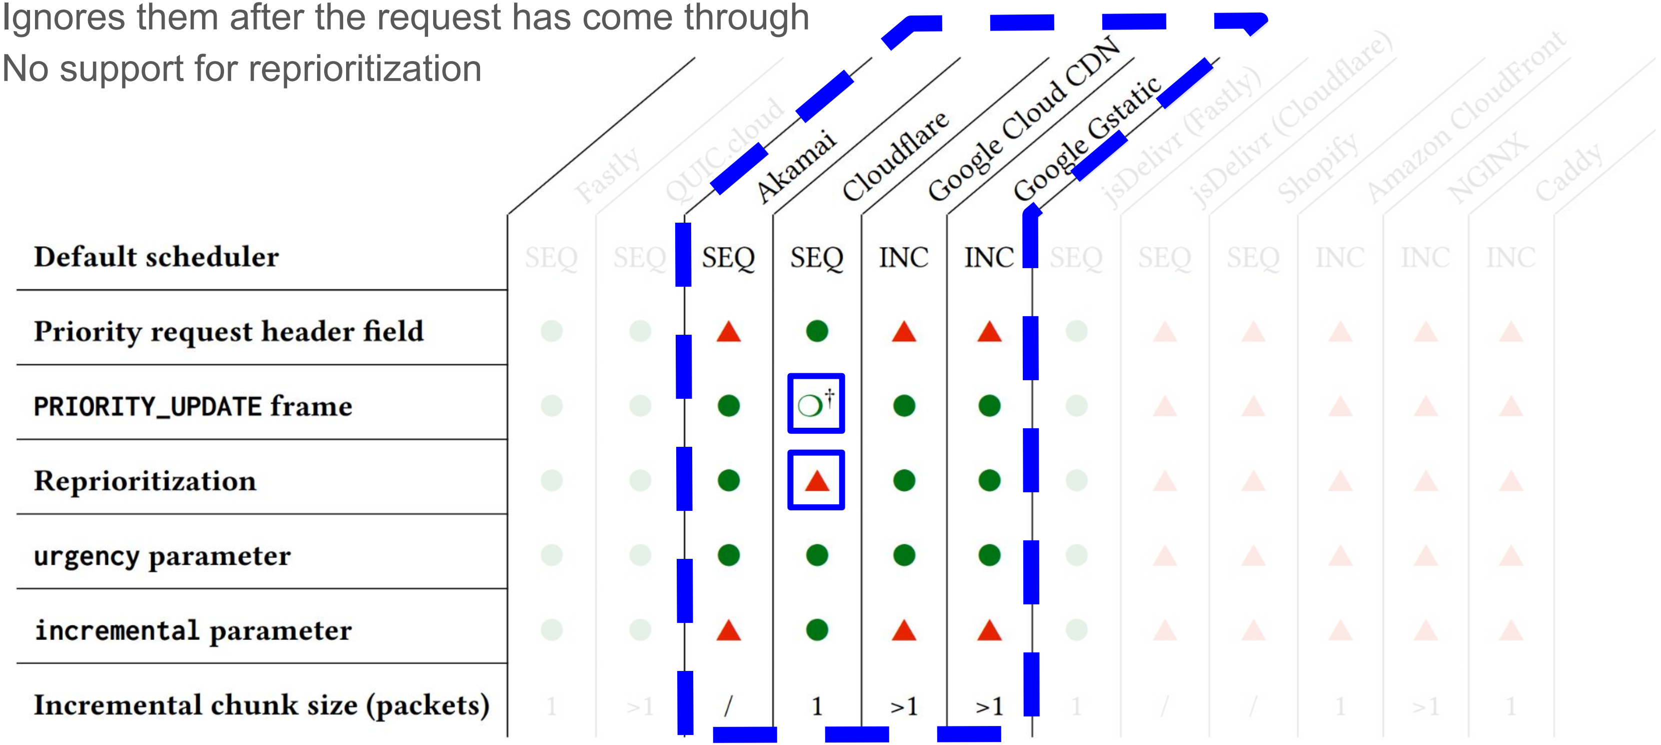

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
`PRIORITY_UPDATE` frame	●	●	●	○†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1

EPS - Servers

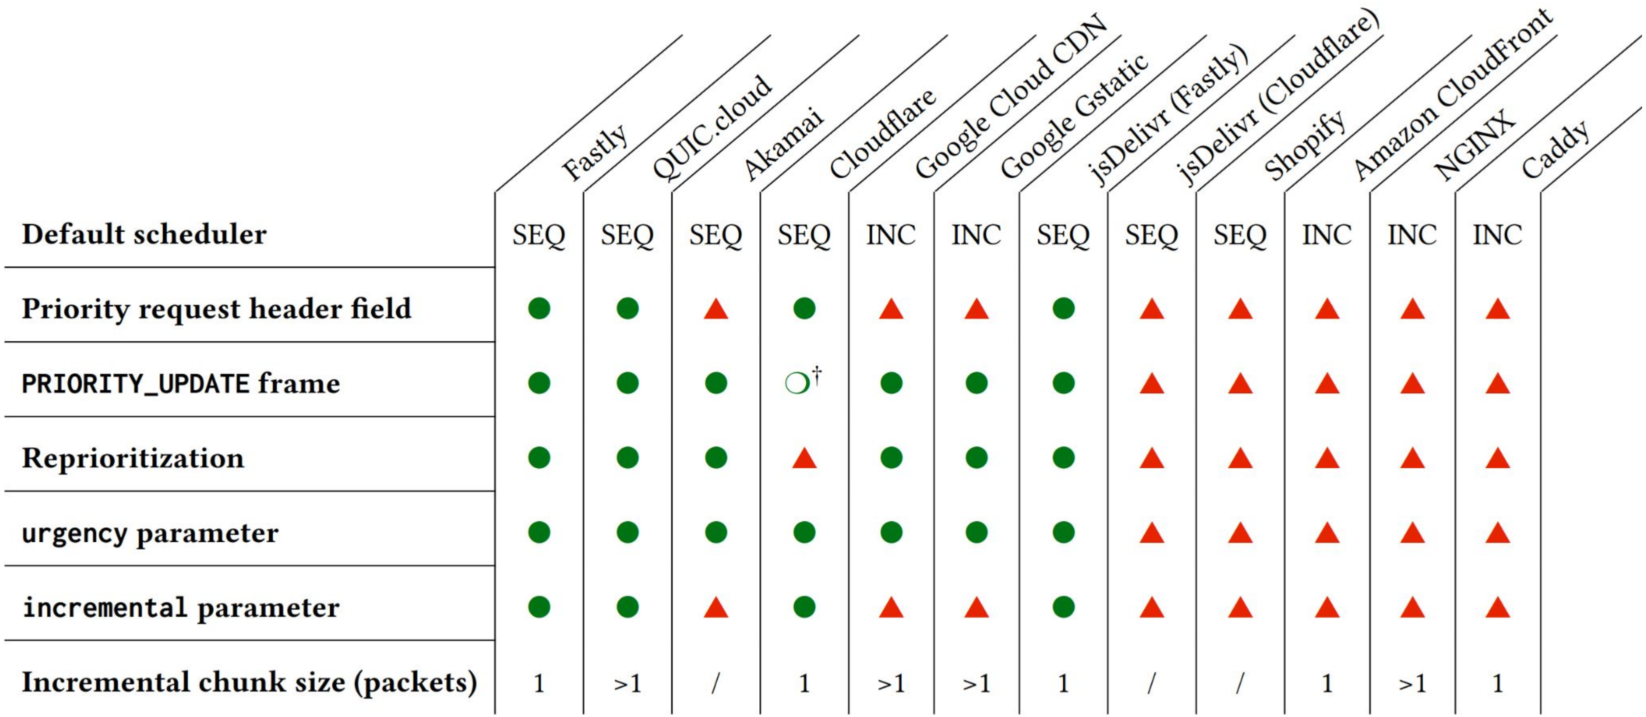

	Fastly	QUIC.cloud	Akamai	Cloudflare	Google Cloud CDN	Google Gstatic	jsDelivr (Fastly)	jsDelivr (Cloudflare)	Shopify	Amazon CloudFront	NGINX	Caddy
Default scheduler	SEQ	SEQ	SEQ	SEQ	INC	INC	SEQ	SEQ	SEQ	INC	INC	INC
Priority request header field	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
PRIORITY_UPDATE frame	●	●	●	❍†	●	●	●	▲	▲	▲	▲	▲
Reprioritization	●	●	●	▲	●	●	●	▲	▲	▲	▲	▲
urgency parameter	●	●	●	●	●	●	●	▲	▲	▲	▲	▲
incremental parameter	●	●	▲	●	▲	▲	●	▲	▲	▲	▲	▲
Incremental chunk size (packets)	1	>1	/	1	>1	>1	1	/	/	1	>1	1