API Rate Limiting Guidelines

To ensure fair and stable access to our API, we enforce rate limits on all endpoints. Please review and follow these guidelines to avoid interruption and optimize integration.

Rate Limit Policy

- Each user is allotted a fixed number of API requests per time window.
- Limits are enforced per API key, per IP, or per account, as applicable.
- Exceeding limits will result in temporary request blocking (HTTP 429).

Standard Limits

API Endpoint	Limit	Time Window
/api/v1/users	1000 requests	1 hour
/api/v1/orders	500 requests	1 hour
All others	300 requests	1 hour

Response Headers

Rate limit details are returned in each response:

- `X-RateLimit-Limit`: Maximum requests per window
- `X-RateLimit-Remaining`: Requests left in current window
- `X-RateLimit-Reset`: Epoch timestamp when the window resets

Sample Headers:

```
X-RateLimit-Limit: 1000
X-RateLimit-Remaining: 58
X-RateLimit-Reset: 1718129966
```

Handling Rate Limit Errors

If you exceed your allocated quota, you will receive a `429 Too Many Requests` response.

```
HTTP/1.1 429 Too Many Requests
Content-Type: application/json
```

```
{
  "error": "Rate limit exceeded. Please retry after 30 seconds."
}
```

Best Practices

- Monitor rate limit headers in all API responses.
- Implement client-side backoff and retry logic.
- Prioritize essential requests to avoid quota exhaustion.
- Contact support for higher limits or bulk use cases.