

PrefectHQ / prefect Public

<> Code Issues 818 Pull requests 39 Discussions Actions Security and

# Commit 7c70ac5



devin-ai-integration[bot] and desertaxle authored 2 weeks ago · ✓ 94 / 99 · Verified



Fix DNS rebinding TOCTOU bypass in validate\_restricted\_url (#21591)

Co-authored-by: Devin AI <158243242+devin-ai-integration[bot]@users.noreply.github.com>

Co-authored-by: Alexander Streed <alex.s@prefect.io>

Co-authored-by: alex.s <ajstreed1@gmail.com>

main (#21591) · prefect-redis-0.2.11 ··· 3.6.28.dev2

1 parent 3458636 commit 7c70ac5

4 files changed

+640 -21

↑ Top ⚙️

Filter files...

src/prefect

blocks

notifications.py

webhook.py

utilities

urls.py

tests/utilities

test\_urls.py

Search within code

src/prefect/blocks/notifications.py

@@ -13,7 +13,11 @@

13 13 from prefect.logging import LogEavesdropper

14	14	<code>from prefect.types import SecretDict</code>
15	15	<code>from prefect.utilities.templating import apply_values, find_placeholders</code>
16		<code>- from prefect.utilities.urls import validate_restricted_url</code>
16		<code>+ from prefect.utilities.urls import (</code>
17		<code>+     SSRFProtectedAsyncHTTPTransport,</code>
18		<code>+     SSRFProtectedHTTPTransport,</code>
19		<code>+     validate_restricted_url,</code>
20		<code>+ )</code>
17	21	
18	22	<code>PREFECT_NOTIFY_TYPE_DEFAULT = "info" # Use a valid apprise type as default</code>
19	23	
		<code>@@ -987,13 +991,18 @@ async def anotify(self, body: str, subject: str   None = None) -&gt; None:</code>
987	991	<code>import httpx</code>
988	992	
989	993	<code>request_args = self._build_request_args(body, subject)</code>
994		<code>+ client_kwargs: dict[str, Any] = {</code>
995		<code>+     "headers": {"user-agent": "Prefect Notifications"},</code>
996		<code>+ }</code>
990	997	<code>if not self.allow_private_urls:</code>
991	998	<code>    validate_restricted_url(request_args["url"])</code>
999		<code>+ # Re-validate at connection time and pin the resolved IP to close</code>
1000		<code>+ # the DNS-rebinding TOCTOU window.</code>
1001		<code>+ client_kwargs["transport"] = SSRFProtectedAsyncHTTPTransport()</code>
992	1002	<code>cookies = request_args.pop("cookies", dict())</code>
1003		<code>+ client_kwargs["cookies"] = cookies</code>
993	1004	<code># make request with httpx</code>
994		<code>- async with httpx.AsyncClient(</code>
995		<code>-     headers={"user-agent": "Prefect Notifications"}, cookies=cookies</code>
996		<code>- ) as client:</code>
1005		<code>+ async with httpx.AsyncClient(**client_kwargs) as client:</code>
997	1006	<code>    resp = await client.request(**request_args)</code>
998	1007	<code>    resp.raise_for_status()</code>
999	1008	
		<code>@@ -1002,13 +1011,16 @@ def notify(self, body: str, subject: str   None = None) -&gt; None:</code>
1002	1011	<code>import httpx</code>
1003	1012	
1004	1013	<code>request_args = self._build_request_args(body, subject)</code>
1014		<code>+ client_kwargs: dict[str, Any] = {</code>

1015	+	"headers": {"user-agent": "Prefect Notifications"},
1016	+	}
1005	1017	if not self.allow_private_urls:
1006	1018	validate_restricted_url(request_args["url"])
1019	+	client_kwargs["transport"] = SSRFProtectedHTTPTransport()
1007	1020	cookies = request_args.pop("cookies", dict())
1021	+	client_kwargs["cookies"] = cookies
1008	1022	# make request with httpx
1009	-	with httpx.Client(
1010	-	headers={"user-agent": "Prefect Notifications"}, cookies=cookies
1011	-	) as client:
1023	+	with httpx.Client(**client_kwargs) as client:
1012	1024	resp = client.request(**request_args)
1013	1025	resp.raise_for_status()
1014	1026	
		↓

▼ src/prefect/blocks/webhook.py		...
↑		@@ -8,12 +8,21 @@
8	8	
9	9	from prefect.blocks.core import Block
10	10	from prefect.types import SecretDict
11	-	from prefect.utilities.urls import validate_restricted_url
11	+	from prefect.utilities.urls import (
12	+	SSRFProtectedAsyncHTTPTransport,
13	+	validate_restricted_url,
14	+	)
12	15	
13	16	# Use a global HTTP transport to maintain a process-wide connection pool for
14	17	# interservice requests
15	18	_http_transport = AsyncHTTPTransport()
16	19	_insecure_http_transport = AsyncHTTPTransport(verify=False)
20	+	# Separate pools for calls that must be protected from DNS-rebinding SSRF. The
21	+	# protected transport validates the resolved IP at connection time and connects
22	+	# to the pre-resolved address, closing the TOCTOU window exploited by DNS
23	+	# rebinding attacks.
24	+	_safe_http_transport = SSRFProtectedAsyncHTTPTransport()
25	+	_safe_insecure_http_transport = SSRFProtectedAsyncHTTPTransport(verify=False)
17	26	
18	27	

```

19 28 class Webhook(Block):
    @@ -53,10 +62,13 @@ class Webhook(Block):
    )
53 62
54 63
55 64 def block_initialization(self) -> None:
56 -     if self.verify:
57 -         self._client = AsyncClient(transport=_http_transport)
65 +     if self.allow_private_urls:
66 +         transport = _http_transport if self.verify else
           _insecure_http_transport
58 67     else:
59 -         self._client = AsyncClient(transport=_insecure_http_transport)
68 +         transport = (
69 +             _safe_http_transport if self.verify else
           _safe_insecure_http_transport
70 +         )
71 +         self._client = AsyncClient(transport=transport)
60 72
61 73     async def call(self, payload: dict[str, Any] | str | None = None) ->
           Response:
62 74         """
    ↓

```

```

src/prefect/utilities/urls.py
... @@ -1,13 +1,18 @@
1 1 import inspect
2 2 import ipaddress
3 3 import socket
4 + import time
4 5 import urllib.parse
6 + from collections.abc import Iterable
5 7 from logging import Logger
6 8 from string import Formatter
7 9 from typing import TYPE_CHECKING, Any, Literal, Optional, Union, cast
8 10 from urllib.parse import urlparse
9 11 from uuid import UUID
10 12
13 + import anyio.to_thread
14 + import httpcore

```

```

15 + import httpx
11 16 from pydantic import BaseModel
12 17
13 18 from prefect import settings
@@ -71,7 +76,14 @@ def validate_restricted_url(url: str) -> None:
71 76     Validate that the provided URL is safe for outbound requests. This
    prevents
72 77     attacks like SSRF (Server Side Request Forgery), where an attacker can make
73 78     requests to internal services (like the GCP metadata service, localhost
    addresses,
74 -     or in-cluster Kubernetes services)
79 +     or in-cluster Kubernetes services).
80 +
81 +     This is a pre-flight check that validates every address the hostname
    resolves
82 +     to via `getaddrinfo`. Because DNS can change between this check and the
83 +     actual HTTP connection, callers that need hardened SSRF protection should
84 +     also use `SSRFProtectedAsyncHTTPTransport` / `SSRFProtectedHTTPTransport`,
85 +     which re-validate at connection time and connect to the pre-resolved IP to
86 +     close the TOCTOU window exploited by DNS rebinding attacks.
75 87
76 88     Args:
77 89         url: The URL to validate.
@@ -100,17 +112,227 @@ def validate_restricted_url(url: str) -> None:
100 112         raise ValueError(f"{url!r} is not a valid URL.")
101 113
102 114     try:
103 -         ip_address = socket.gethostbyname(hostname)
104 -         ip = ipaddress.ip_address(ip_address)
115 +         _validate_resolved_hostname(hostname)
116 +     except _RestrictedHostError as exc:
117 +         raise ValueError(f"{url!r} is not a valid URL. {exc}")
118 +
119 +
120 + class _RestrictedHostError(Exception):
121 +     """Internal exception raised when a hostname resolves to a restricted
    address."""
122 +

```

```
123 +
124 + def _validate_resolved_hostname(hostname: str) -> list[str]:
125 +     """Resolve `hostname` and validate every returned address.
126 +
127 +     Returns the list of resolved IPs (as strings) in resolution order. Raises
128 +     `_RestrictedHostError` if any resolved address is private, or if the
129 +     hostname cannot be resolved.
130 +
131 +     Using `getaddrinfo` (rather than `gethostbyname`, which returns only the
132 +     first A record) closes an SSRF bypass where a hostname publishes both a
133 +     public and a private A/AAAA record: every resolved address is checked.
134 +     """
135 +     # IP literal: validate directly.
136 +     try:
137 +         ip = ipaddress.ip_address(hostname)
138 +     except ValueError:
139 +         pass
140 +     else:
141 +         if ip.is_private:
142 +             raise _RestrictedHostError(f"It resolves to the private address
143 + {ip}.")
143 +         return [str(ip)]
144 +
145 +     try:
146 +         addrinfos = socket.getaddrinfo(hostname, None, type=socket.SOCK_STREAM)
105 147         except socket.gaierror:
148 +             raise _RestrictedHostError("It could not be resolved.")
149 +
150 +         resolved: list[str] = []
151 +         for addrinfo in addrinfos:
152 +             sockaddr = addrinfo[4]
153 +             ip_str = sockaddr[0]
154 +             # Strip IPv6 zone identifier if present (e.g. "fe80::1%eth0")
155 +             ip_str = ip_str.split("%", 1)[0]
106 156         try:
107 -             ip = ipaddress.ip_address(hostname)
157 +             ip = ipaddress.ip_address(ip_str)
108 158         except ValueError:
109 -             raise ValueError(f"{url!r} is not a valid URL. It could not be
110 +             resolved.")
```

```
159 +         continue
160 +         if ip.is_private:
161 +             raise _RestrictedHostError(f"It resolves to the private address
{ip}.")
162 +         if ip_str not in resolved:
163 +             resolved.append(ip_str)
110 164
111 -         if ip.is_private:
112 -             raise ValueError(
113 -                 f"{url!r} is not a valid URL. It resolves to the private address
{ip}."
165 +         if not resolved:
166 +             raise _RestrictedHostError("It could not be resolved.")
167 +
168 +         return resolved
169 +
170 +
171 + class _SSRFProtectedAsyncBackend(httplib.AsyncNetworkBackend):
172 +     """An `httplib.AsyncNetworkBackend` that validates resolved addresses.
173 +
174 +     Wraps an existing backend and, on each `connect_tcp` call, resolves the
175 +     hostname itself, rejects any resolved address that is private, and then
176 +     connects to the validated IP directly (rather than the hostname) so that
177 +     the underlying backend cannot re-resolve to a different address.
178 +
179 +     TLS SNI / certificate validation is unaffected because httplib passes the
180 +     original hostname to `start_tls` via `server_hostname`, independently of
181 +     the host used for the TCP connection.
182 +     """
183 +
184 +     def __init__(self, wrapped: httplib.AsyncNetworkBackend) -> None:
185 +         self._wrapped = wrapped
186 +
187 +     async def connect_tcp(
188 +         self,
189 +         host: str,
190 +         port: int,
191 +         timeout: Optional[float] = None,
192 +         local_address: Optional[str] = None,
193 +         socket_options: Optional[Iterable[Any]] = None,
```

```
194 +     ) -> httpcore.AsyncNetworkStream:
195 +         # Resolve in a worker thread so the event loop is not blocked during
196 +         # DNS lookups (which can be slow or intermittently failing).
197 +         validated_ips = await anyio.to_thread.run_sync(
198 +             _resolve_and_validate_for_connect, host
199 +         )
200 +         last_exc: Optional[BaseException] = None
201 +         deadline = time.monotonic() + timeout if timeout is not None else None
202 +         for ip in validated_ips:
203 +             remaining = _remaining_timeout(deadline)
204 +             if remaining == 0.0:
205 +                 break
206 +             try:
207 +                 return await self._wrapped.connect_tcp(
208 +                     ip,
209 +                     port,
210 +                     timeout=remaining,
211 +                     local_address=local_address,
212 +                     socket_options=socket_options,
213 +                 )
214 +             except (httpcore.ConnectError, httpcore.ConnectTimeout, OSError) as
exc:
215 +                 last_exc = exc
216 +                 assert last_exc is not None
217 +                 raise last_exc
218 +
219 +     async def connect_unix_socket(
220 +         self,
221 +         path: str,
222 +         timeout: Optional[float] = None,
223 +         socket_options: Optional[Iterable[Any]] = None,
224 +     ) -> httpcore.AsyncNetworkStream:
225 +         return await self._wrapped.connect_unix_socket(
226 +             path, timeout=timeout, socket_options=socket_options
227 +         )
228 +
229 +     async def sleep(self, seconds: float) -> None:
230 +         await self._wrapped.sleep(seconds)
231 +
232 +
```

```
233 + class _SSRFProtectedSyncBackend(httpcore.NetworkBackend):
234 +     """Synchronous counterpart of `_SSRFProtectedAsyncBackend`."""
235 +
236 +     def __init__(self, wrapped: httpcore.NetworkBackend) -> None:
237 +         self._wrapped = wrapped
238 +
239 +     def connect_tcp(
240 +         self,
241 +         host: str,
242 +         port: int,
243 +         timeout: Optional[float] = None,
244 +         local_address: Optional[str] = None,
245 +         socket_options: Optional[Iterable[Any]] = None,
246 +     ) -> httpcore.NetworkStream:
247 +         validated_ips = _resolve_and_validate_for_connect(host)
248 +         last_exc: Optional[BaseException] = None
249 +         deadline = time.monotonic() + timeout if timeout is not None else None
250 +         for ip in validated_ips:
251 +             remaining = _remaining_timeout(deadline)
252 +             if remaining == 0.0:
253 +                 break
254 +             try:
255 +                 return self._wrapped.connect_tcp(
256 +                     ip,
257 +                     port,
258 +                     timeout=remaining,
259 +                     local_address=local_address,
260 +                     socket_options=socket_options,
261 +                 )
262 +             except (httpcore.ConnectError, httpcore.ConnectTimeout, OSError) as exc:
263 +                 last_exc = exc
264 +                 assert last_exc is not None
265 +                 raise last_exc
266 +
267 +     def connect_unix_socket(
268 +         self,
269 +         path: str,
270 +         timeout: Optional[float] = None,
271 +         socket_options: Optional[Iterable[Any]] = None,
```

```
272 +     ) -> httpcore.NetworkStream:
273 +         return self._wrapped.connect_unix_socket(
274 +             path, timeout=timeout, socket_options=socket_options
275 +         )
276 +
277 +     def sleep(self, seconds: float) -> None:
278 +         self._wrapped.sleep(seconds)
279 +
280 +
281 +     def _remaining_timeout(deadline: Optional[float]) -> Optional[float]:
282 +         """Return the time left until `deadline`, or `None` if no deadline is set.
283 +
284 +         Clamps to `0.0` when the deadline has already passed. Callers use this to
285 +         share a single connect-timeout budget across multiple per-IP retries so
286 +         that connect time stays bounded (roughly) by the caller's timeout rather
287 +         than scaling with the number of resolved addresses.
288 +         """
289 +         if deadline is None:
290 +             return None
291 +         return max(0.0, deadline - time.monotonic())
292 +
293 +
294 +     def _resolve_and_validate_for_connect(host: str) -> list[str]:
295 +         """Resolve `host` and return all safe IPs to connect to.
296 +
297 +         Every returned IP has been validated against the private-address blocklist;
298 +         callers iterate them in order and retry on connect failures so that dual-
299 +         stack hostnames still work in single-stack environments.
300 +
301 +         Raises `httpcore.ConnectError` if any resolved address is private or if the
302 +         hostname cannot be resolved. The returned IPs are passed to the underlying
303 +         network backend as IP literals, so it will not perform further DNS
304 +         resolution – eliminating the DNS rebinding TOCTOU window.
305 +         """
306 +         try:
307 +             return _validate_resolved_hostname(host)
308 +         except _RestrictedHostError as exc:
309 +             raise httpcore.ConnectError(f"Refusing to connect to {host!r}: {exc}")
310 +     from None
```

```

311 +
312 + class SSRFProtectedAsyncHTTPTransport(httpx.AsyncHTTPTransport):
313 +     """An `httpx.AsyncHTTPTransport` that guards against DNS rebinding SSRF.
314 +
315 +     Behaves identically to `httpx.AsyncHTTPTransport` except that, for every
316 +     request, the hostname is resolved, every resolved address is checked
317 +     against the private-address blocklist, and the connection is made to the
318 +     specific validated IP. This closes the TOCTOU window between a pre-flight
319 +     `validate_restricted_url` check and the actual HTTP connection.
320 +     """
321 +
322 +     def __init__(self, *args: Any, **kwargs: Any) -> None:
323 +         super().__init__(*args, **kwargs)
324 +         self._pool._network_backend = _SSRFProtectedAsyncBackend(
325 +             self._pool._network_backend
326 +         )
327 +
328 +
329 + class SSRFProtectedHTTPTransport(httpx.HTTPTransport):
330 +     """Synchronous counterpart of `SSRFProtectedAsyncHTTPTransport`."""
331 +
332 +     def __init__(self, *args: Any, **kwargs: Any) -> None:
333 +         super().__init__(*args, **kwargs)
334 +         self._pool._network_backend = _SSRFProtectedSyncBackend(
335 +             self._pool._network_backend
114 336         )
115 337
116 338

```



tests/utilities/test\_urls.py



@@ -1,8 +1,13 @@

```

1 1 import concurrent.futures
2 + import socket
3 + import threading
2 4 import uuid
3 5 from datetime import timedelta
4 - from typing import Any, Literal
6 + from typing import Any, Literal, cast
7 + from unittest.mock import patch

```

```

5      8
9      + import httpcore
10     + import httpx
6      11     import pytest
7      12
8      13     from prefect.blocks.webhook import Webhook
@@ -14,7 +19,14 @@
14     19     from prefect.settings import PREFECT_API_URL, PREFECT_UI_URL,
        temporary_settings
15     20     from prefect.states import StateType
16     21     from prefect.types._datetime import now
17     - from prefect.utilities.urls import url_for, validate_restricted_url
22     + from prefect.utilities.urls import (
23     +     SSRFProtectedAsyncHTTPTransport,
24     +     SSRFProtectedHTTPTransport,
25     +     _SSRFProtectedAsyncBackend,
26     +     _SSRFProtectedSyncBackend,
27     +     url_for,
28     +     validate_restricted_url,
29     + )
18     30     from prefect.variables import Variable
19     31
20     32     MOCK_PREFECT_UI_URL = "https://ui.prefect.io"
@@ -459,3 +471,364 @@ def
@@ -471,3 +471,364 @@ def
test_url_for_with_additional_format_kwargs_raises_if_placeholder_not_replace
459   471         match="Unable to generate URL for worker because the following keys are
        missing: work_pool_name",
460   472     ):
461   473         url_for(obj="worker", obj_id="123e4567-e89b-12d3-a456-42661417400")
474   +
475   +
476   + def _fake_getaddrinfo(ips: list[str]):
477   +     """Return a `getaddrinfo` stub that yields `ips` as resolved addresses."""
478   +
479   +     def _impl(host: str, *_args: Any, **_kwargs: Any):
480   +         results = []
481   +         for ip in ips:
482   +             try:
483   +                 family = socket.AF_INET6 if ":" in ip else socket.AF_INET
484   +                 except Exception:

```

```
485 +         family = socket.AF_INET
486 +         results.append((family, socket.SOCK_STREAM, 0, "", (ip, 0)))
487 +     return results
488 +
489 +     return _impl
490 +
491 +
492 + class TestValidateRestrictedUrlMultiRecord:
493 +     def test_rejects_when_any_resolved_ip_is_private(self):
494 +         """A hostname that resolves to both a public and a private address
495 +         must be rejected – `gethostbyname` returned only the first record and
496 +         could be bypassed by rotating records."""
497 +         with patch(
498 +             "prefect.utilities.urls.socket.getaddrinfo",
499 +             side_effect=_fake_getaddrinfo(["8.8.8.8", "10.0.0.1"]),
500 +         ):
501 +             with pytest.raises(ValueError, match="private address 10.0.0.1"):
502 +                 validate_restricted_url("https://example.com")
503 +
504 +     def test_accepts_when_all_resolved_ips_are_public(self):
505 +         with patch(
506 +             "prefect.utilities.urls.socket.getaddrinfo",
507 +             side_effect=_fake_getaddrinfo(["8.8.8.8", "1.1.1.1"]),
508 +         ):
509 +             validate_restricted_url("https://example.com")
510 +
511 +     def test_rejects_ipv6_mapped_private(self):
512 +         """A hostname that only resolves to an IPv6 address in a private
513 +         range is rejected."""
514 +         with patch(
515 +             "prefect.utilities.urls.socket.getaddrinfo",
516 +             side_effect=_fake_getaddrinfo(["fc00::1"]),
517 +         ):
518 +             with pytest.raises(ValueError, match="private address"):
519 +                 validate_restricted_url("https://example.com")
520 +
521 +
522 + class TestSSRFProtectedTransportTOCTOU:
523 +     """The protected transport must reject private IPs at connection time,
524 +     even when an initial validation via `validate_restricted_url` succeeded.
```

```
525 +     This simulates the DNS rebinding attack described in BOUNTY-422."""
526 +
527 +     async def test_async_transport_rejects_rebinding_to_private_ip(self):
528 +         transport = SSRFProtectedAsyncHTTPTransport()
529 +         # At connection time, DNS resolves to a private IP.
530 +         with patch(
531 +             "prefect.utilities.urls.socket.getaddrinfo",
532 +             side_effect=_fake_getaddrinfo(["127.0.0.1"]),
533 +         ):
534 +             async with httpx.AsyncClient(transport=transport) as client:
535 +                 with pytest.raises(
536 +                     httpx.ConnectError, match="private address 127.0.0.1"
537 +                 ):
538 +                     await client.get("http://example.com")
539 +
540 +     def test_sync_transport_rejects_rebinding_to_private_ip(self):
541 +         transport = SSRFProtectedHTTPTransport()
542 +         with patch(
543 +             "prefect.utilities.urls.socket.getaddrinfo",
544 +             side_effect=_fake_getaddrinfo(["192.168.1.1"]),
545 +         ):
546 +             with httpx.Client(transport=transport) as client:
547 +                 with pytest.raises(
548 +                     httpx.ConnectError, match="private address 192.168.1.1"
549 +                 ):
550 +                     client.get("http://example.com")
551 +
552 +     async def test_async_transport_rejects_multi_record_with_private(self):
553 +         """Even if one resolved A record is public, a sibling private record
554 +         must fail the validation (defends against DNS responses that mix
555 +         public and private addresses)."""
556 +         transport = SSRFProtectedAsyncHTTPTransport()
557 +         with patch(
558 +             "prefect.utilities.urls.socket.getaddrinfo",
559 +             side_effect=_fake_getaddrinfo(["8.8.8.8", "10.0.0.1"]),
560 +         ):
561 +             async with httpx.AsyncClient(transport=transport) as client:
562 +                 with pytest.raises(
563 +                     httpx.ConnectError, match="private address 10.0.0.1"
564 +                 ):
```

```
565 +         await client.get("http://example.com")
566 +
567 +     async def test_async_transport_rejects_unresolvable_host(self):
568 +         transport = SSRFProtectedAsyncHTTPTransport()
569 +         with patch(
570 +             "prefect.utilities.urls.socket.getaddrinfo",
571 +             side_effect=socket.gaierror("resolution failed"),
572 +         ):
573 +             async with httpx.AsyncClient(transport=transport) as client:
574 +                 with pytest.raises(httpx.ConnectError, match="could not be
resolved"):
575 +                     await client.get("http://example.com")
576 +
577 +     async def test_async_transport_rejects_private_ip_literal(self):
578 +         transport = SSRFProtectedAsyncHTTPTransport()
579 +         async with httpx.AsyncClient(transport=transport) as client:
580 +             with pytest.raises(httpx.ConnectError, match="private address"):
581 +                 await client.get("http://127.0.0.1")
582 +
583 +     def test_transport_wraps_existing_network_backend(self):
584 +         """Ensure the transport replaces the pool's network backend with the
585 +         SSRF-protected wrapper while keeping the underlying backend
586 +         available for actual connections."""
587 +         atransport = SSRFProtectedAsyncHTTPTransport()
588 +         assert isinstance(atransport._pool._network_backend,
_SSRFProtectedAsyncBackend)
589 +
590 +         stransport = SSRFProtectedHTTPTransport()
591 +         assert isinstance(stransport._pool._network_backend,
_SSRFProtectedSyncBackend)
592 +
593 +
594 +     class TestSSRFProtectedBackendPinsIP:
595 +         """The protected backends must connect to the resolved IP (a literal),
596 +         not the original hostname, so that the underlying backend cannot
597 +         re-resolve DNS and pick up a different (private) address."""
598 +
599 +         async def test_async_backend_connects_to_validated_ip(self):
600 +             class RecordingBackend(httpcore.AsyncNetworkBackend):
601 +                 def __init__(self):
```

```
602 +         self.connected_host: str | None = None
603 +
604 +         async def connect_tcp(
605 +             self, host, port, timeout=None, local_address=None,
606 +             socket_options=None
607 +         ):
608 +             self.connected_host = host
609 +             # Return a mock stream – we don't actually connect.
610 +             raise httpcore.ConnectError("stop here")
611 +
612 +         async def connect_unix_socket(self, *args, **kwargs):
613 +             raise NotImplementedError
614 +
615 +         async def sleep(self, seconds):
616 +             pass
617 +
618 +         inner = RecordingBackend()
619 +         backend = _SSRFProtectedAsyncBackend(inner)
620 +
621 +         with patch(
622 +             "prefect.utilities.urls.socket.getaddrinfo",
623 +             side_effect=_fake_getaddrinfo(["8.8.8.8"]),
624 +         ):
625 +             with pytest.raises(httpcore.ConnectError, match="stop here"):
626 +                 await backend.connect_tcp("example.com", 443)
627 +
628 +             assert inner.connected_host == "8.8.8.8"
629 +
630 +     class TestSSRFProtectedBackendFallbackAcrossIPs:
631 +         """When `getaddrinfo` returns several safe IPs (e.g. a dual-stack AAAA+A
632 +         record set), the protected backends must try each validated IP in order
633 +         and only raise once every one has failed. Pinning a single IP would
634 +         break dual-stack hostnames in single-stack environments."""
635 +
636 +         async def test_async_backend_falls_back_to_next_validated_ip(self):
637 +             class FlakyBackend(httpcore.AsyncNetworkBackend):
638 +                 def __init__(self):
639 +                     self.connected_hosts: list[str] = []
640 +
```

```
641 +         async def connect_tcp(
642 +             self, host, port, timeout=None, local_address=None,
643 +             socket_options=None
644 +         ):
645 +             self.connected_hosts.append(host)
646 +             # First IP is unreachable, second succeeds.
647 +             if len(self.connected_hosts) == 1:
648 +                 raise httpcore.ConnectError("network unreachable")
649 +             return cast(httpcore.AsyncNetworkStream, object())
650 +
651 +         async def connect_unix_socket(self, *args, **kwargs):
652 +             raise NotImplementedError
653 +
654 +         async def sleep(self, seconds):
655 +             pass
656 +
657 +         inner = FlakyBackend()
658 +         backend = _SSRFProtectedAsyncBackend(inner)
659 +
660 +         with patch(
661 +             "prefect.utilities.urls.socket.getaddrinfo",
662 +             side_effect=_fake_getaddrinfo(["2606:4700:4700::1111",
663 +             "93.184.216.34"]),
664 +         ):
665 +             await backend.connect_tcp("example.com", 443)
666 +
667 +             assert inner.connected_hosts == ["2606:4700:4700::1111",
668 +             "93.184.216.34"]
669 +
670 +         async def test_async_backend_raises_last_error_when_all_fail(self):
671 +             class AlwaysFailingBackend(httpcore.AsyncNetworkBackend):
672 +                 def __init__(self):
673 +                     self.attempts = 0
674 +
675 +                 async def connect_tcp(
676 +                     self, host, port, timeout=None, local_address=None,
677 +                     socket_options=None
678 +                 ):
679 +                     self.attempts += 1
680 +                     raise httpcore.ConnectError(f"attempt {self.attempts} failed")
```

```
677 +
678 +         async def connect_unix_socket(self, *args, **kwargs):
679 +             raise NotImplementedError
680 +
681 +         async def sleep(self, seconds):
682 +             pass
683 +
684 +         inner = AlwaysFailingBackend()
685 +         backend = _SSRFProtectedAsyncBackend(inner)
686 +
687 +         with patch(
688 +             "prefect.utilities.urls.socket.getaddrinfo",
689 +             side_effect=_fake_getaddrinfo(["2606:4700:4700::1111",
690 +             "93.184.216.34"]),
691 +         ):
692 +             with pytest.raises(httpcore.ConnectError, match="attempt 2
693 +             failed"):
694 +                 await backend.connect_tcp("example.com", 443)
695 +
696 +             assert inner.attempts == 2
697 +
698 +         def test_sync_backend_falls_back_to_next_validated_ip(self):
699 +             class FlakyBackend(httpcore.NetworkBackend):
700 +                 def __init__(self):
701 +                     self.connected_hosts: list[str] = []
702 +
703 +                 def connect_tcp(
704 +                     self, host, port, timeout=None, local_address=None,
705 +                     socket_options=None
706 +                 ):
707 +                     self.connected_hosts.append(host)
708 +                     if len(self.connected_hosts) == 1:
709 +                         raise httpcore.ConnectError("network unreachable")
710 +                     return cast(httpcore.NetworkStream, object())
711 +
712 +                 def connect_unix_socket(self, *args, **kwargs):
713 +                     raise NotImplementedError
714 +
715 +                 def sleep(self, seconds):
716 +                     pass
```

```
714 +
715 +     inner = FlakyBackend()
716 +     backend = _SSRFProtectedSyncBackend(inner)
717 +
718 +     with patch(
719 +         "prefect.utilities.urls.socket.getaddrinfo",
720 +         side_effect=_fake_getaddrinfo(["2606:4700:4700::1111",
721 + "93.184.216.34"]),
722 +     ):
723 +         backend.connect_tcp("example.com", 443)
724 +         assert inner.connected_hosts == ["2606:4700:4700::1111",
725 + "93.184.216.34"]
726 +
727 + class TestSSRFProtectedBackendTimeoutBudget:
728 +     """The total connect timeout across all per-IP retries must stay bounded
729 +     by the caller's `timeout`. Without a shared budget, a hostname with N
730 +     resolved addresses could take up to N * timeout to fail."""
731 +
732 +     async def test_async_backend_shares_timeout_budget_across_retries(self):
733 +         class SlowFailingBackend(httpcore.AsyncNetworkBackend):
734 +             def __init__(self):
735 +                 self.timeouts: list[float | None] = []
736 +
737 +             async def connect_tcp(
738 +                 self, host, port, timeout=None, local_address=None,
739 +                 socket_options=None
740 +             ):
741 +                 self.timeouts.append(timeout)
742 +                 raise httpcore.ConnectError("nope")
743 +
744 +             async def connect_unix_socket(self, *args, **kwargs):
745 +                 raise NotImplementedError
746 +
747 +             async def sleep(self, seconds):
748 +                 pass
749 +
750 +         inner = SlowFailingBackend()
751 +         backend = _SSRFProtectedAsyncBackend(inner)
```

```
751 +
752 +     with patch(
753 +         "prefect.utilities.urls.socket.getaddrinfo",
754 +         side_effect=_fake_getaddrinfo(["2606:4700:4700::1111",
755 + "93.184.216.34"]),
756 +     ):
757 +         with pytest.raises(httpcore.ConnectError):
758 +             await backend.connect_tcp("example.com", 443, timeout=5.0)
759 +
760 +         # Each attempt receives a timeout no greater than the caller's budget;
761 +         # retries inherit the *remaining* budget, not a fresh one.
762 +         assert len(inner.timeouts) == 2
763 +         for t in inner.timeouts:
764 +             assert t is not None
765 +             assert t <= 5.0
766 +             assert inner.timeouts[1] <= inner.timeouts[0]
767 +
768 +     def test_sync_backend_shares_timeout_budget_across_retries(self):
769 +         class SlowFailingBackend(httpcore.NetworkBackend):
770 +             def __init__(self):
771 +                 self.timeouts: list[float | None] = []
772 +
773 +             def connect_tcp(
774 +                 self, host, port, timeout=None, local_address=None,
775 +                 socket_options=None
776 +             ):
777 +                 self.timeouts.append(timeout)
778 +                 raise httpcore.ConnectError("nope")
779 +
780 +             def connect_unix_socket(self, *args, **kwargs):
781 +                 raise NotImplementedError
782 +
783 +             def sleep(self, seconds):
784 +                 pass
785 +
786 +         inner = SlowFailingBackend()
787 +         backend = _SSRFProtectedSyncBackend(inner)
788 +
789 +         with patch(
790 +             "prefect.utilities.urls.socket.getaddrinfo",
```

```
789 +         side_effect=_fake_getaddrinfo(["2606:4700:4700::1111",
790 +         "93.184.216.34"]),
791 +     ):
792 +         with pytest.raises(httpcore.ConnectError):
793 +             backend.connect_tcp("example.com", 443, timeout=5.0)
794 +
795 +         assert len(inner.timeouts) == 2
796 +         for t in inner.timeouts:
797 +             assert t is not None
798 +             assert t <= 5.0
799 +             assert inner.timeouts[1] <= inner.timeouts[0]
800 +
801 + class TestSSRFProtectedAsyncBackendNonBlockingDNS:
802 +     """`_SSRFProtectedAsyncBackend.connect_tcp` must not block the event loop
803 +     during DNS resolution; it runs `getaddrinfo` in a worker thread."""
804 +
805 +     async def test_async_backend_runs_getaddrinfo_off_the_event_loop(self):
806 +         loop_thread = threading.get_ident()
807 +         observed_threads: list[int] = []
808 +
809 +         def fake_getaddrinfo(*args, **kwargs):
810 +             observed_threads.append(threading.get_ident())
811 +             return [(socket.AF_INET, socket.SOCK_STREAM, 0, "",
812 + ("93.184.216.34", 0))]
813 +
814 +         class StopBackend(httpcore.AsyncNetworkBackend):
815 +             async def connect_tcp(self, *args, **kwargs):
816 +                 raise httpcore.ConnectError("stop")
817 +
818 +             async def connect_unix_socket(self, *args, **kwargs):
819 +                 raise NotImplementedError
820 +
821 +             async def sleep(self, seconds):
822 +                 pass
823 +
824 +         backend = _SSRFProtectedAsyncBackend(StopBackend())
825 +
826 +         with patch(
```

```
826 +         "prefect.utilities.urls.socket.getaddrinfo",
      side_effect=fake_getaddrinfo
827 +     ):
828 +         with pytest.raises(httpcore.ConnectError, match="stop"):
829 +             await backend.connect_tcp("example.com", 443)
830 +
831 +         assert observed_threads, "getaddrinfo should have been invoked"
832 +         assert observed_threads[0] != loop_thread, (
833 +             "getaddrinfo must run off the event loop thread"
834 +         )
```

## Comments 0



Please [sign in](#) to comment.