

inter.link

Inter.link's automation journey

Κώστας Δρόγγος

costas@inter.link

GRNOG 17 - 06.12.2024



WE PROVIDE **SUSTAINABLE**
CONNECTIVITY OF THE
HIGHEST **QUALITY** THROUGH
AUTOMATION & INNOVATION

Is that about ansible again?

Promise: I'm only going to mention Ansible **once**

So what is this about?

Is that about ansible again?

Promise: I'm only going to mention Ansible **once**

So what is this about?

"As a Customer I want to order a service at midnight, patch cables right after and use it immediately"

Challenges

"As a Customer I want to order a service at midnight, patch cables right after and use it immediately"

- Transparent pricing
- Portal should reflect state of infrastructure
- Portal should allow input by the Customer (e.g. vlan id or bgp password)
- Transparent pricing → transparent Billing
- Services must be deployed safely to the common infrastructure
- Maintenance windows might be happening in some PoP at the same time
- Other jobs might be running at the same time



Find your desired location







Search an address or city to find nearby datacenters...

Your existing ports

-  Vienna 10Gbps
SID764
- SCO4134 (Port)
 AT.VIE.VIE2.F00.200.R02.B07.U12.X09/10
-  Interxion Vienna VIE1/VIE2
 Louis-Haefliger-Gasse 10, 1210 Vienna

Inter.link locations

Locations with 100/400GE PoP owned and operated by Inter.link. Services are available for immediate provisioning.

- | | | |
|---|--|--|
|  Interxion Vienna VIE1/VIE2
VIE1-AT
Vienna
Louis-Haefliger-Gasse 10 |  NTT Vienna 1 Data Center (VIE1)
VIE2-AT
Vienna
Computerstraße 4 |  CE Colo Prague
PRG1-CZ
Prague
Nad Elektrarnou 1428/47 |
|  Digital Realty Copenhagen (CPH1, CPH2)
CPH1-DK
København
Industriparken 20A |  Telehouse - Paris 2 (Voltaire - Léon Frot) incl. Paris 1/3 via Campus XC
PARI-FR
Paris
65, rue Léon Frot |  Datacenter One HAM1 Ahrensburg
AHR1-DE
Ahrensburg
Carl-Backhaus-Str. 21c |

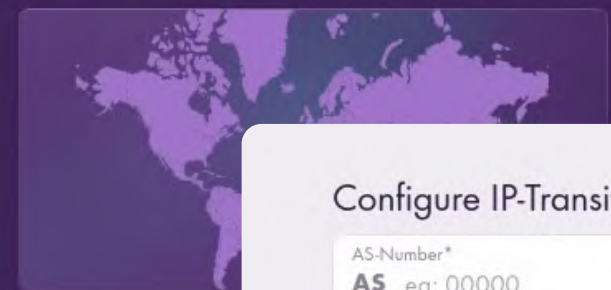
Select contract term

1M 12M **24 MONTHS** 36M

IP Transit

Service Summary

Location



Contract term

Configure IP-Transit

AS-Number* AS eg: 00000	AS-SET eg: AS-INTERDEMO
Prefix IPv4 Limit* eg: 20	Prefix IPv6 Limit* eg: 20
Session Password	Route Advertisement* Full Routing Table

Sync BGP prefix limits automatically ⓘ

Configure IP-Transit

Purchase Reference
eg: PO-1

Committed Data Rate

- | | | | |
|--------------------------------|--------------------------------|---------------------------------|--------------------------------|
| 1 Gbps
350,00€ MRC | 2 Gbps
600,00€ MRC | 3 Gbps
816,67€ MRC | 4 Gbps
1.033,33€ MRC |
| 6 Gbps
1.380,00€ MRC | 8 Gbps
1.640,00€ MRC | 10 Gbps
1.900,00€ MRC | |

Aggregated Billing


1M 12M **24 MONTHS** 36M

IP Transit

Service Summary

Location

Vienna
NTT Vienna 1 Data Center (VIE1)



Port Details

Port Speed 10 Gbps

One Time Fee 250,00€

VLAN Untagged

BGP Session

AS-SET ASundefined ⓘ

Route Advertisement Full Routing Table

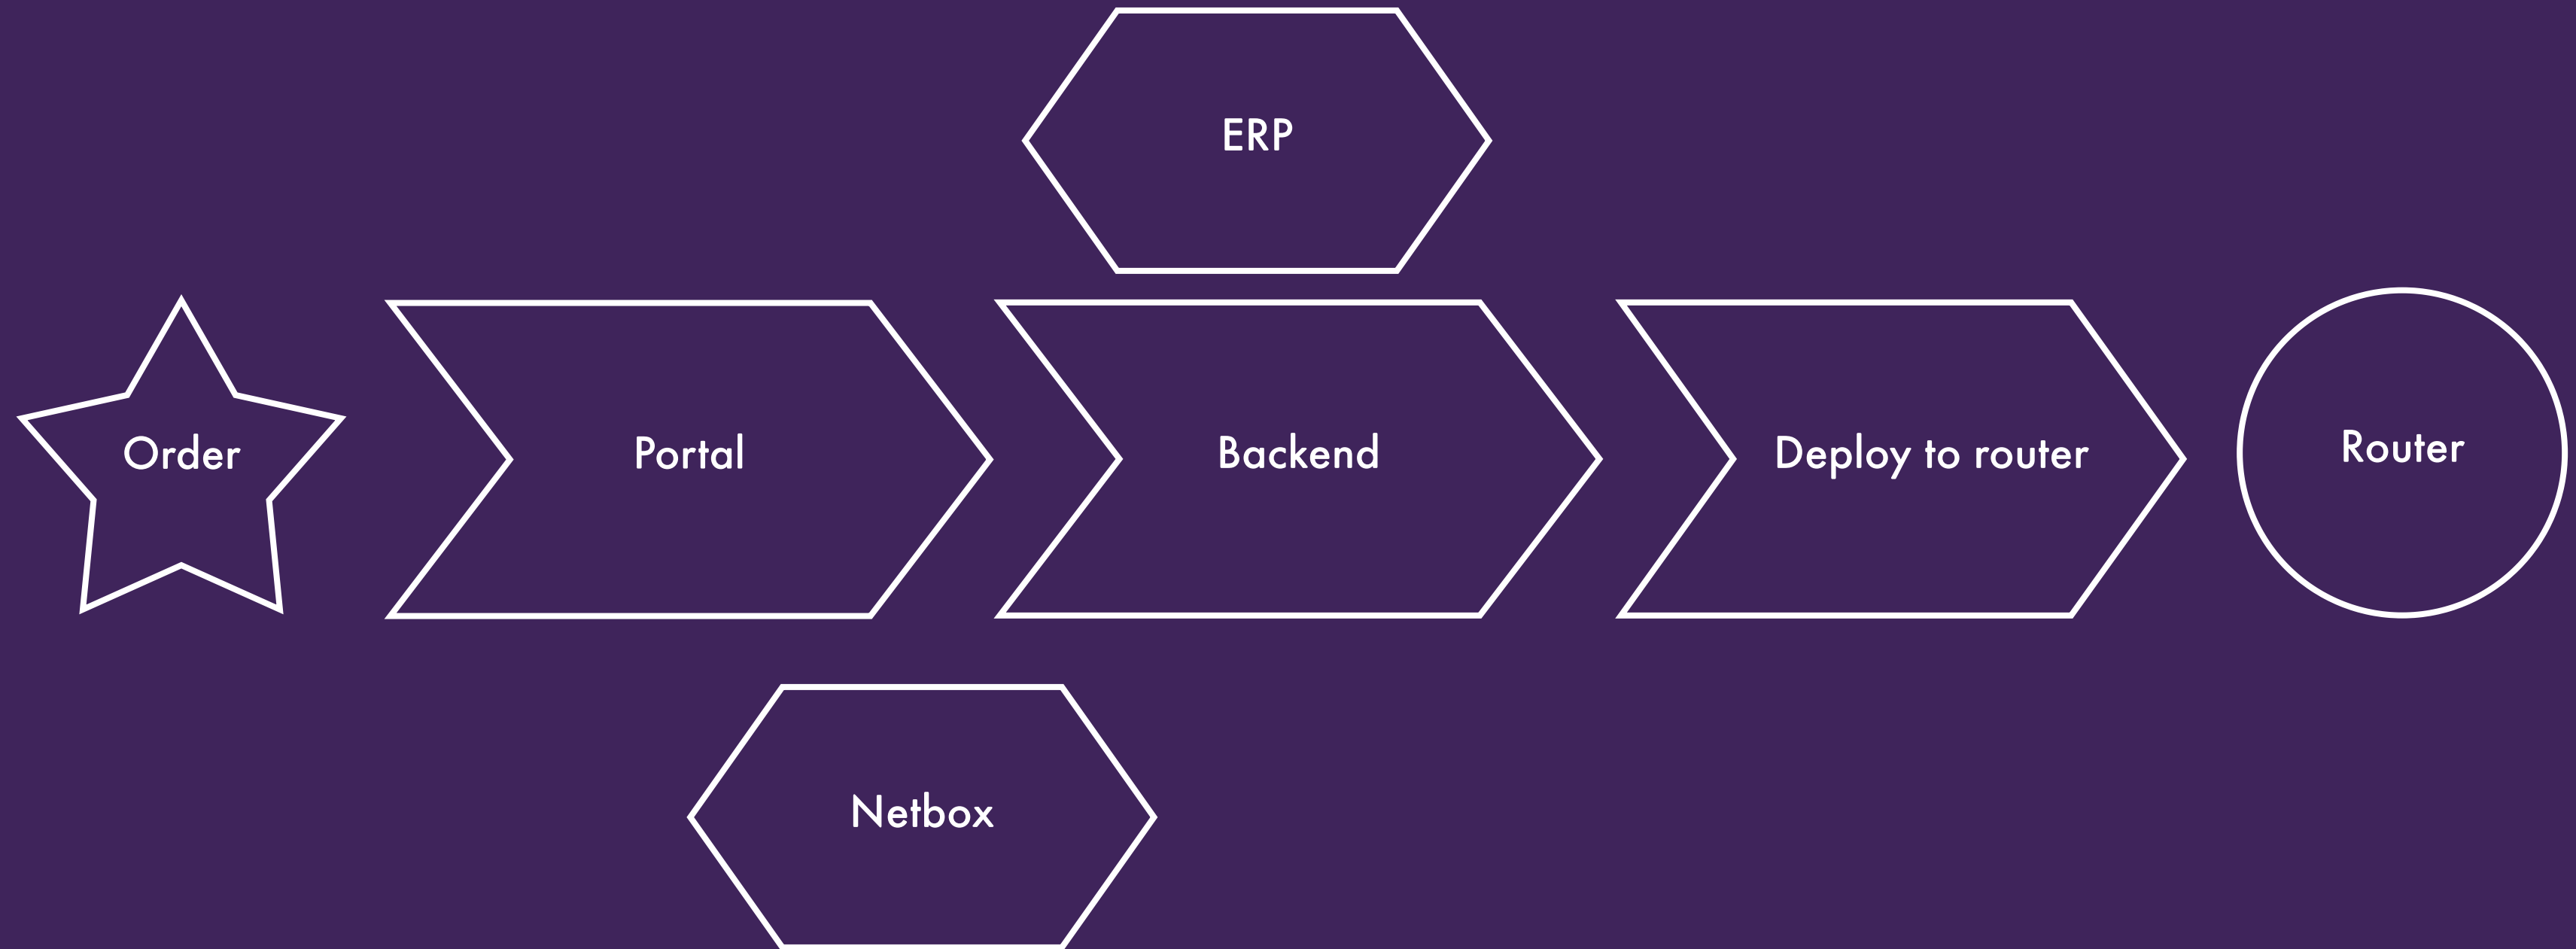
Prefixes

IPv4 /31 2 Addresses 0,00€

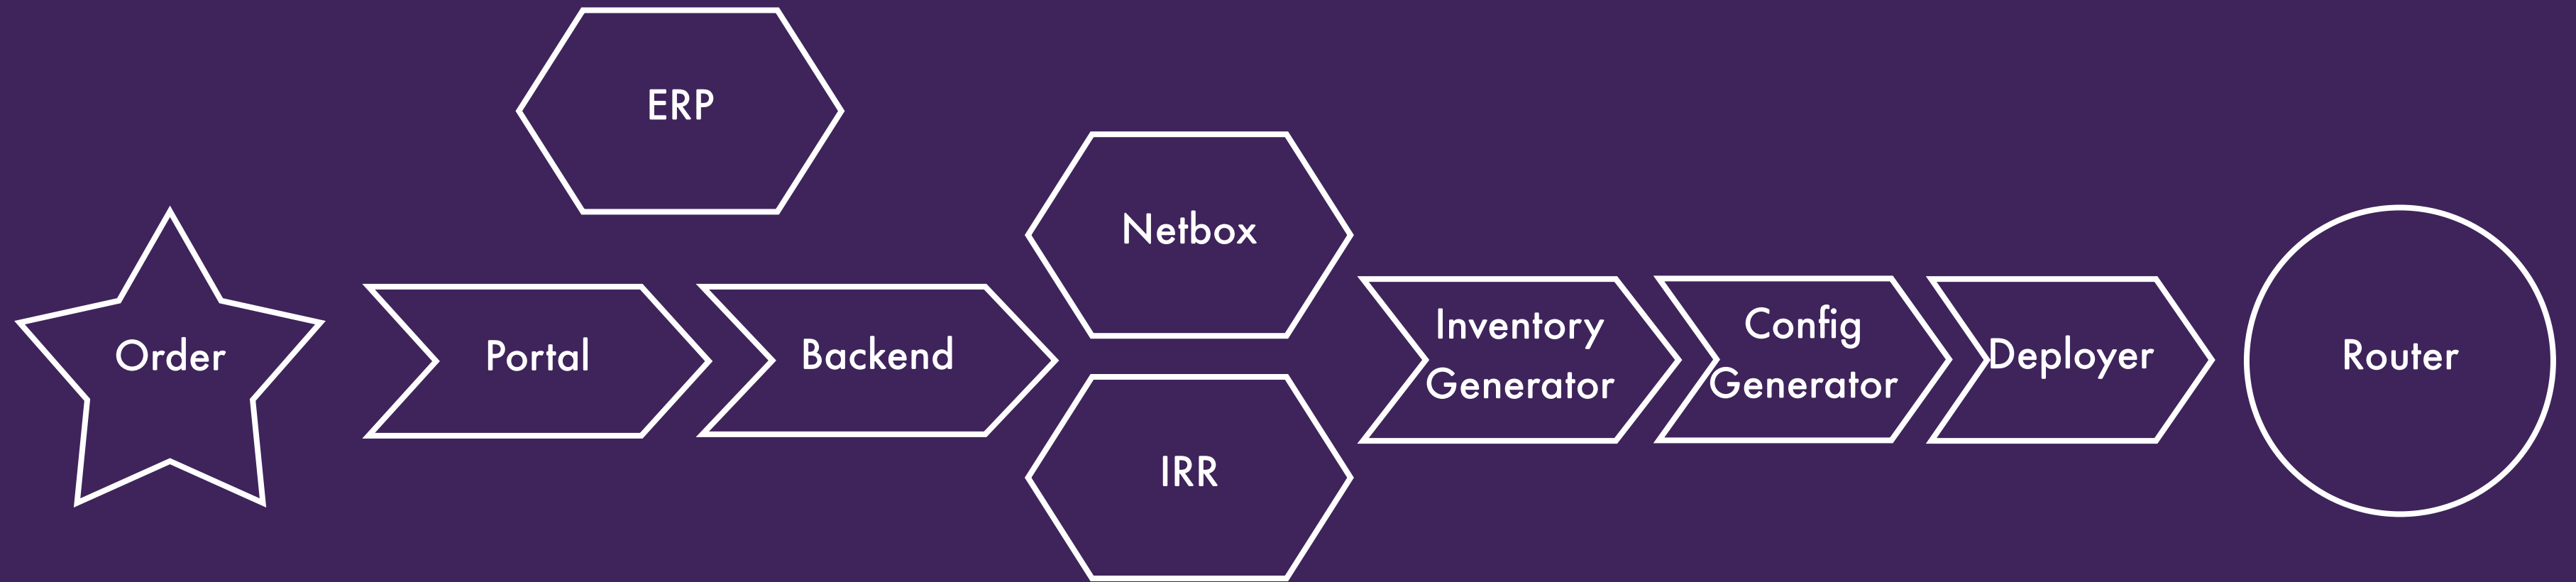
IPv6 /127 2 Addresses 0,00€

Contract term 24 months

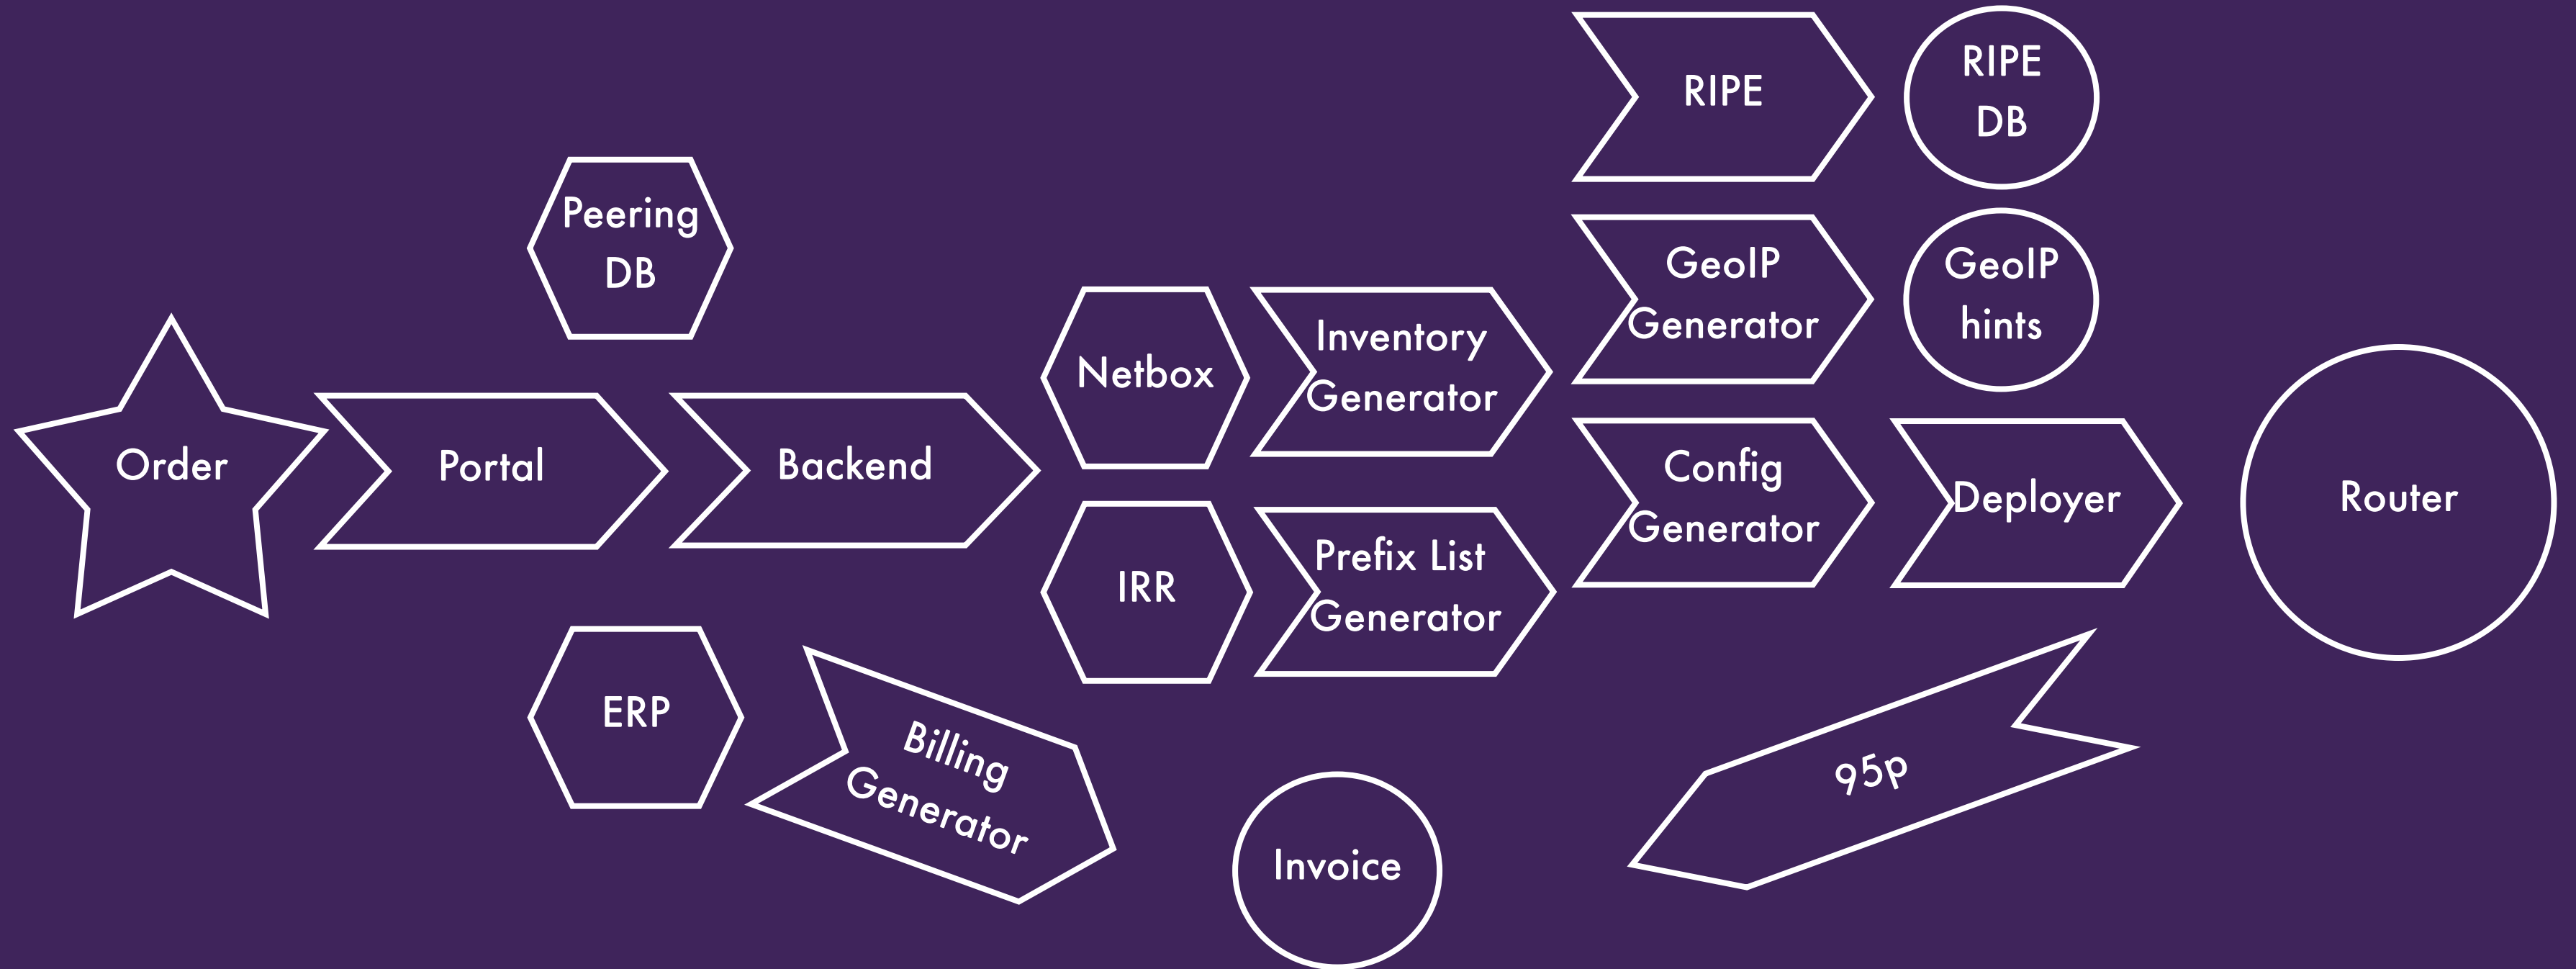
Order an IP-T



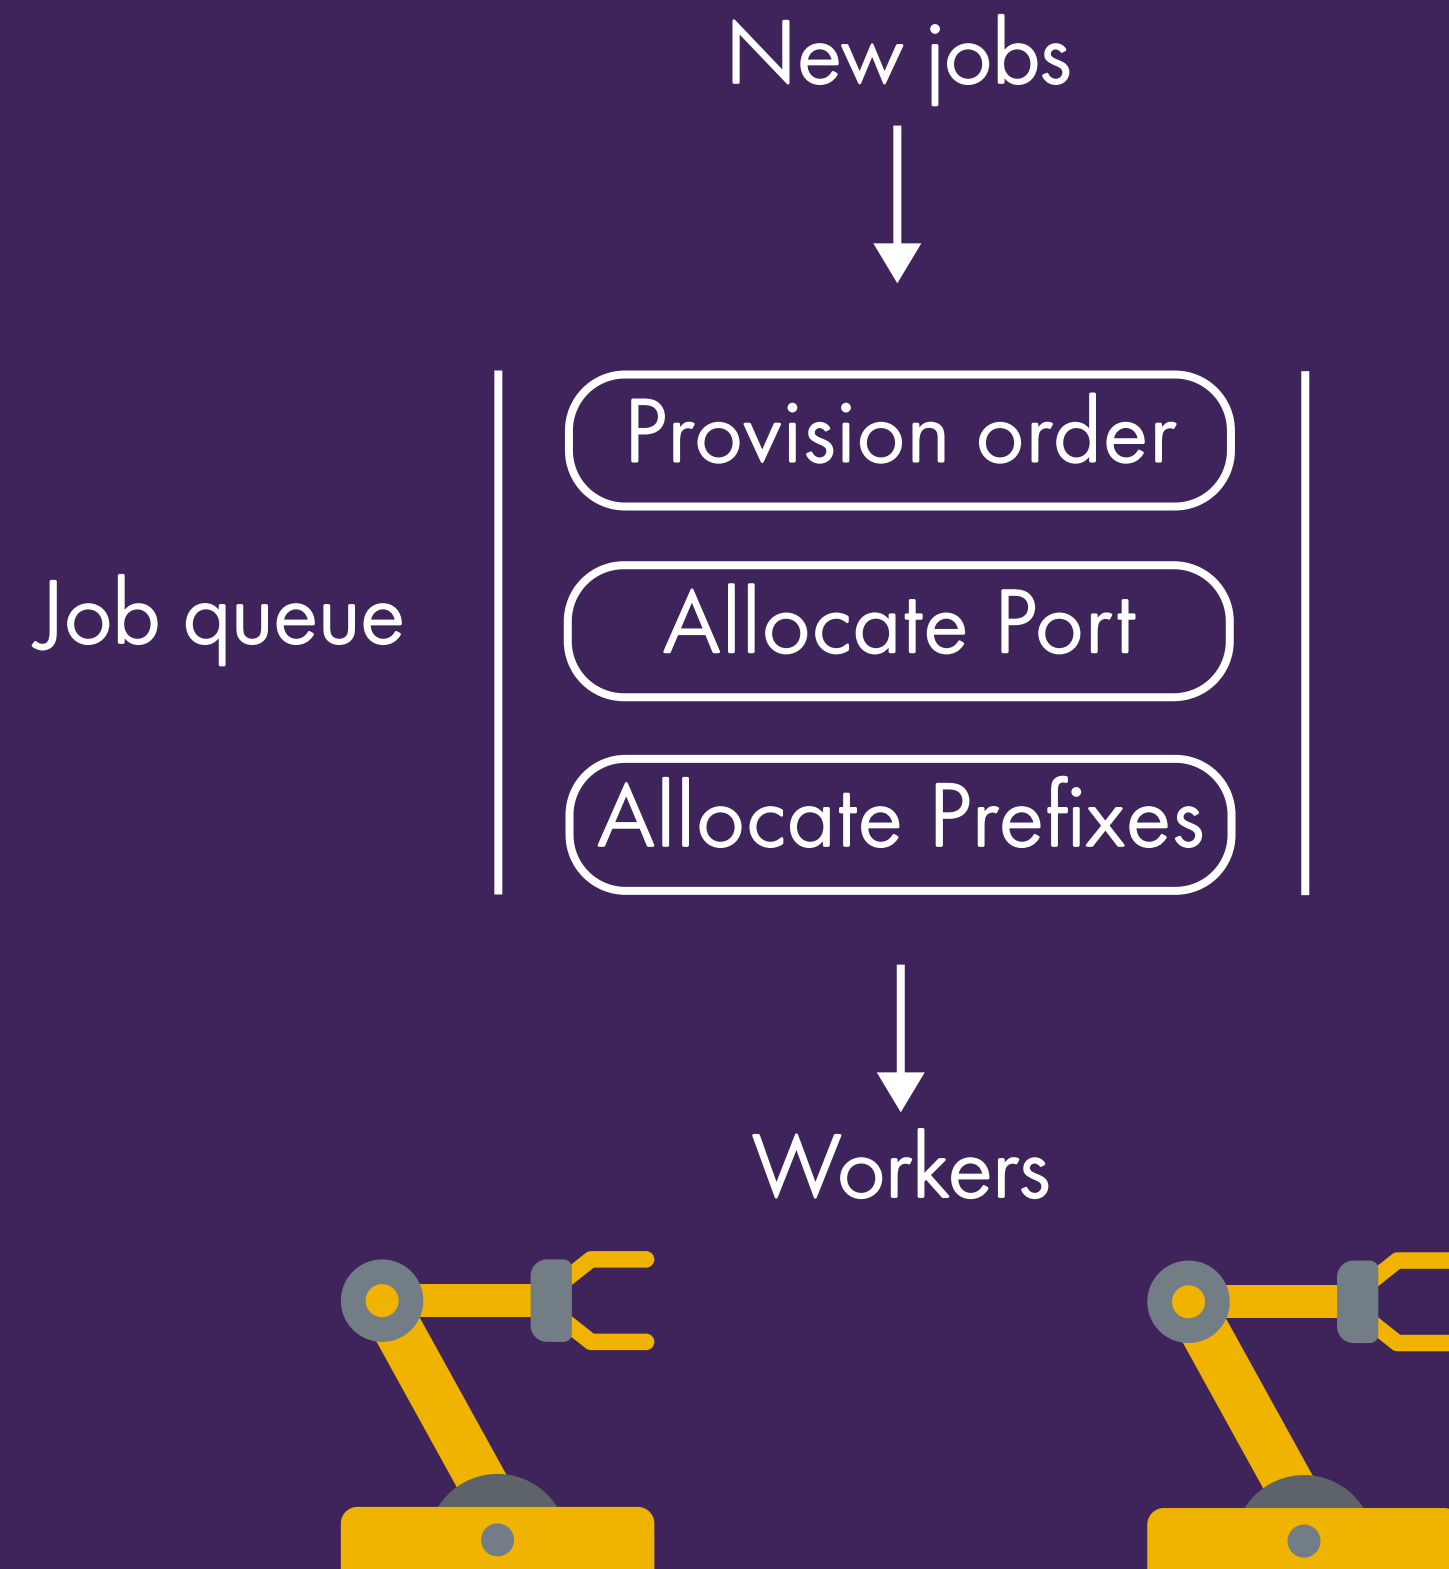
Order an IP-T



Order an IP-T



Order an IP-T



Input validation

How many ways can an AS number be entered?

- 1234
- AS1234
- ASN1234

What about an AS-SET?

- AS-TEST
- AS1234:AS-TEST
- RIPE::AS1234:AS-TEST
- RIPE::AS-TEST

Configure IP-Transit

AS-Number*
AS asdas

AS-SET
eg: AS-INTERDEMO

Prefix IPv4 Limit*
eg: 20

Prefix IPv6 Limit*
eg: 20

Session Password

Route Advertisement*
Full Routing Table

⚠ value is not a valid integer

We need to ensure that all input is validated and normalized

```
46 @classmethod
47 def validate_as_number_public(cls, as_number: int) -> int:
48     if (
49         (as_number == 0) # IANA reserved
50         or (as_number == 23456) # RFC 7607
51         or (as_number >= 64496 and as_number <= 64512) # RFC 4893
52         or (as_number >= 65536 and as_number <= 65552) # RFC 5398
53         or (as_number >= 64512 and as_number <= 65535) # RFC 5398
54         or (
55             as_number >= 4200000000 and as_number <= 4294967296
56         ) # RFC 6996
57         or (as_number == 65535) # RFC 6996
58         or (as_number == 4294967295) # RFC 6996
59         or (as_number >= 65552 and as_number <= 131072) # RFC 6996
60     ):
61         raise ValueError(f"{as_number} is not a public AS number")
62     return as_number
63
```

Service Representation

Billing plans change, services change (what do you mean you need two ports and no IPv4?)

- Break services into basic, reusable components that we can mix and match (and override)

<input type="checkbox"/>	Service Component Name	Type ↓	Product Compo...
1	<input type="checkbox"/> Untagged	Bundle	VLAN
2	<input type="checkbox"/> AT.VIE.VIE2.F00.200.R02.B07.U12.X09...	Bundle	Port
3	<input type="checkbox"/> 45.153.83.66/31	Bundle	Prefix
4	<input type="checkbox"/> Billing Group	Bundle	Billing Group
5	<input type="checkbox"/> Burst	Bundle	Billing Group Burst
6	<input type="checkbox"/> IP Bandwidth 1000 Mbps	Bundle	IP Bandwidth
7	<input type="checkbox"/> 2a11:4140:d002::a/127	Bundle	Prefix
8	<input type="checkbox"/> A test company by Costas	Bundle	BGP Session
9	<input type="checkbox"/> Setup Fee	Bundle	Setup
10	<input type="checkbox"/> A test company by Costas	Bundle	BGP Session

- ... But at the same time, keep a service entity around to reflect the current mix of components

```
896 class IpTransitProduct(BaseProduct):
897     cdr: SpeedMbps
898     components: List[IpTransitProductComponent]
899     product_gid = ProductIds.IP_TRANSIT
900
901     class Config:
902         smart_union = (
903             True # otherwise, it tries confuses vlan.type and port.type
904         )
905
906     _dependencies: List[Tuple[DepKey, DepKey]] = [
907         (DepKey.bgp_v4, DepKey.prefix_v4),
908         (DepKey.bgp_v6, DepKey.prefix_v6),
909         (DepKey.prefix_v4, DepKey.vlan),
910         (DepKey.prefix_v6, DepKey.vlan),
911         (DepKey.vlan, DepKey.port),
912         (DepKey.vlan, DepKey.lag),
913         (DepKey.ports, DepKey.lag),
914     ]
915
```


Idempotency

Executing the same functionality with the same input, shouldn't produce a different output

How to ensure?

Repeated tests! Do I get the same thing every time I push the same arguments?

- unit tests
- provision tests
- end to end tests



Re-Queing

Put an order back into the pipeline?

- Service in the middle of the pipeline crashed
- Bug
- Change in details (say, add a BGP MD5 password)

All that is created (and allocated) until that point shouldn't be recreated → job saves state

How to ensure?

- Idempotency ✓
- Global IDs ✓
- Dependencies between jobs ✗
- Locks ✓



Tight Coupling

What happens when a service is being deployed when the request comes in?

- Queues - retry provisioning when something fails (idempotency!)
- Make all communication asynchronous
 - REST APIs, stateless execution
 - Services connect back when they have results
 - Locks for critical sections
- State of things cached in different layers (inventory)
 - cache busting when connecting back with the result
- Generated configs persisted in Git, state in DBs



Network: Router configuration

partial or **full configs**?

- lock when generating the inventory
- lock when deploying a device
- CI/CD, lint, syntax, diff, tests
- autogenerated descriptions

GET	/api/network/device-lock	Get Device Locks All
POST	/api/network/device-lock	Create Device Lock
DELETE	/api/network/device-lock	Delete Device Lock

✓ Merge request pipeline #1573573582 passed ✓ ✓ ✓ ✓ 📄
Merge request pipeline passed for 1b15a734 4 hours ago ⓘ

Doing all things the "ansible way" stopped working for us

deploy-container before & after moving prefix-list compilation out of ansible (green is the new one), last 7 days

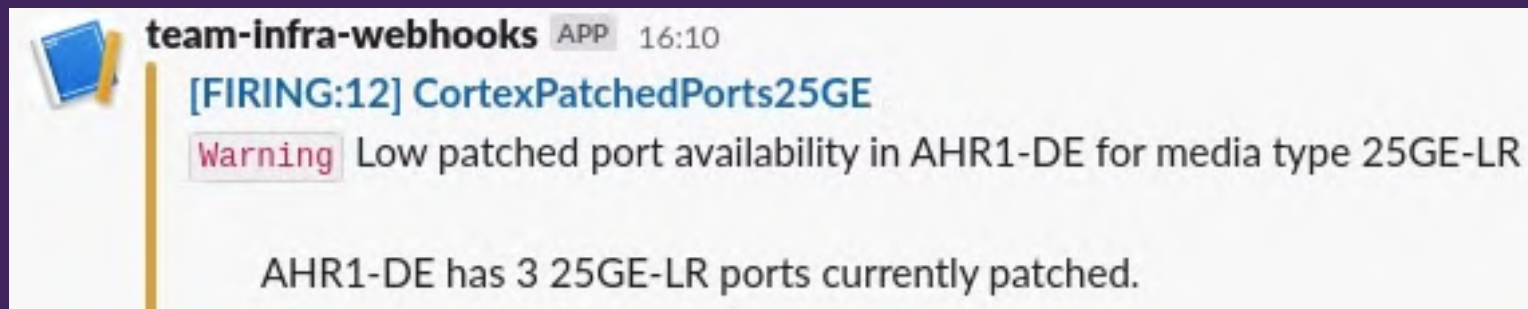
2024-05-31-092907_400x187_screenshot.png ▼

Name	Last *	Max
— deploy-container-d48744c99-46dlw	249 MIB	594 MIB
— deploy-container-d48744c99-fj84s	2.44 GIB	2.44 GIB
— deploy-container-d48744c99-tsrkn	132 MIB	3.28 GIB
— deploy-container-d48744c99-zdxrt	3.65 GIB	5.19 GIB

Network: Cabling

Need to ensure we always have enough ports (or enough time to cable more)

- No loose coupling and caches here → Synchronous check
- Spreadsheets don't play well after the first few hundred ports
- Remote hands take time, shipping optics takes time



A screenshot of a Slack message from the 'team-infra-webhooks' app, timestamped 16:10. The message title is '[FIRING:12] CortexPatchedPorts25GE'. The main content is a warning: 'Warning Low patched port availability in AHR1-DE for media type 25GE-LR'. Below the warning, it states 'AHR1-DE has 3 25GE-LR ports currently patched.'

POST	/api/network/port-allocation	Create Port
DELETE	/api/network/port-allocation	Delete Port
GET	/api/network/port-allocation	Get Port Deprecated
PATCH	/api/network/port-allocation	Update Port





Thank you!
Questions?

NETWORK MAP
OUR GLOBAL SERVICE NETWORK

- America**
- San Jose (SJC)
 - Los Angeles (LAX)
 - Dallas (DAL)
 - Miami (MIA)
 - Atlanta (ATL)
 - Ashburn (ASH)

- Europe**
- | | | | | |
|------------------|------------------|-----------------|-------------------|---------------------------|
| Aarhus (AAR) | Copenhagen (CPH) | Ljubljana (LJU) | Nuremberg (NUE) | Stockholm (STN) |
| Ahrensburg (AHR) | Dublin (DBL) | London (LON) | Oslo (OSL) | Stuttgart (STR) |
| Amsterdam (AMS) | Dusseldorf (DUS) | Madrid (MAD) | Paris (PAR) | Thessaloniki (SKG) |
| Berlin (BER) | Frankfurt (FRA) | Marseille (MRS) | Prague (PRG) | Vienna (VIE) |
| Budapest (BUD) | Hamburg (HAM) | Milan (MIL) | Skanderborg (SKG) | Zurich (ZRH) |
| Bucharest (BUH) | Helsinki (HEL) | Munich (MUC) | Sofia (SOF) | |

- Asia**
- Hong Kong (HKG)
 - Singapore (SIN)
 - Sydney (SYD)
 - Tokyo (TOK)

Q1 2025 