



# ESnet

ENERGY SCIENCES NETWORK

# ESnet's Orchestration Perspective

**Marc Körner**

Software Engineer  
Orchestration and Core Data

Energy Sciences Network (ESnet)  
Lawrence Berkeley National Laboratory

I2 TechEX

Minneapolis, MN  
August 2023

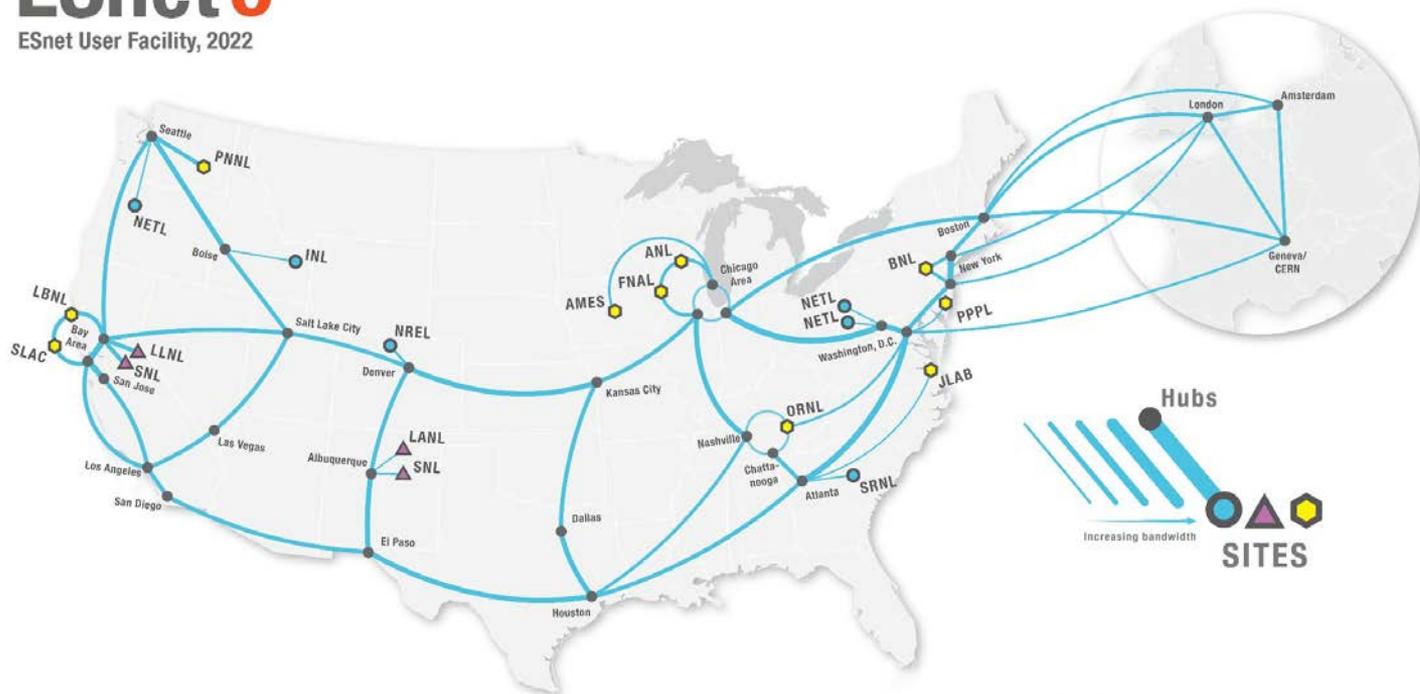


U.S. DEPARTMENT OF  
**ENERGY**  
Office of Science



# ESnet6

ESnet User Facility, 2022



## ESnet in numbers:

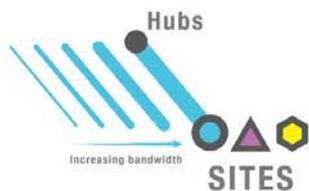
Thousands of miles of fiber cables, including transatlantic cables

380 locations with racks and equipment to track

346 Core links between routers

300 Customer facing Interfaces, 123 of which are 100G

Multi Platform environment with lots of interoperability needs



### Office of Science National Laboratories

- AMES** Ames Laboratory (Ames, IA)
- ANL** Argonne National Laboratory (Argonne, IL)
- BNL** Brookhaven National Laboratory (Upton, NY)
- FNAL** Fermi National Accelerator Laboratory (Batavia, IL)
- JLAB** Thomas Jefferson National Accelerator Facility (Newport News, VA)
- LBL** Lawrence Berkeley National Laboratory (Berkeley, CA)
- LLNL** Lawrence Livermore National Laboratory (Livermore, CA)
- ORNL** Oak Ridge National Laboratory (Oak Ridge, TN)
- PNNL** Pacific Northwest National Laboratory (Richland, WA)
- PPPL** Princeton Plasma Physics Laboratory (Princeton, NJ)
- SLAC** SLAC National Accelerator Laboratory (Menlo Park, CA)

### NNSA Laboratories

- LANL** Los Alamos National Laboratory (Los Alamos, NM)
- LLNL** Lawrence Livermore National Laboratory (Livermore, CA)
- SNL** Sandia National Laboratory (Albuquerque, NM; Livermore, CA)

### Other DOE Laboratories

- INL** Idaho National Laboratory (Idaho Falls, ID)
- NETL** National Energy Technology Laboratory (Morgantown, WV; Pittsburgh, PA; Albany, OR)
- NREL** National Renewable Energy Laboratory (Golden, CO)
- SRNL** Savannah River National Laboratory (Aiken, SC)

# What is orchestration?

Orchestration is defined as the automated arrangement, coordination, and management of computer systems, middleware, and services within the network.

-- Wikipedia



Workflow management implementing a repeatable pattern of steps for user-based interactions within our operations support system;

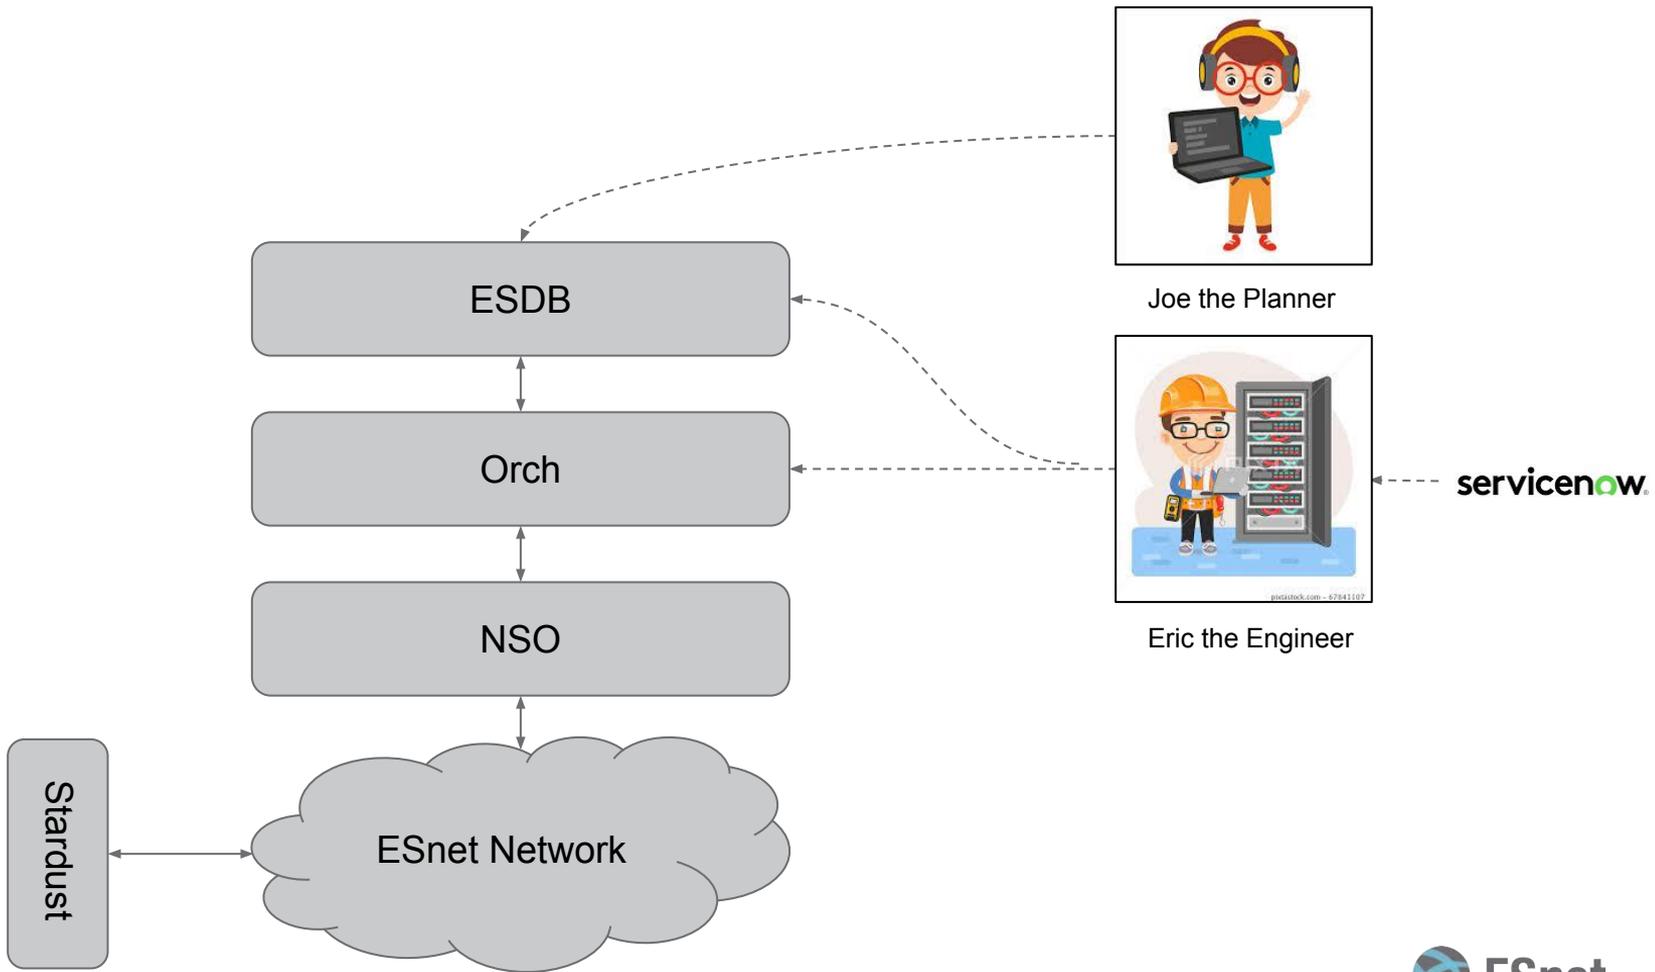
Automated service provisioning to programmatically change configuration across network, compute, and application resources.

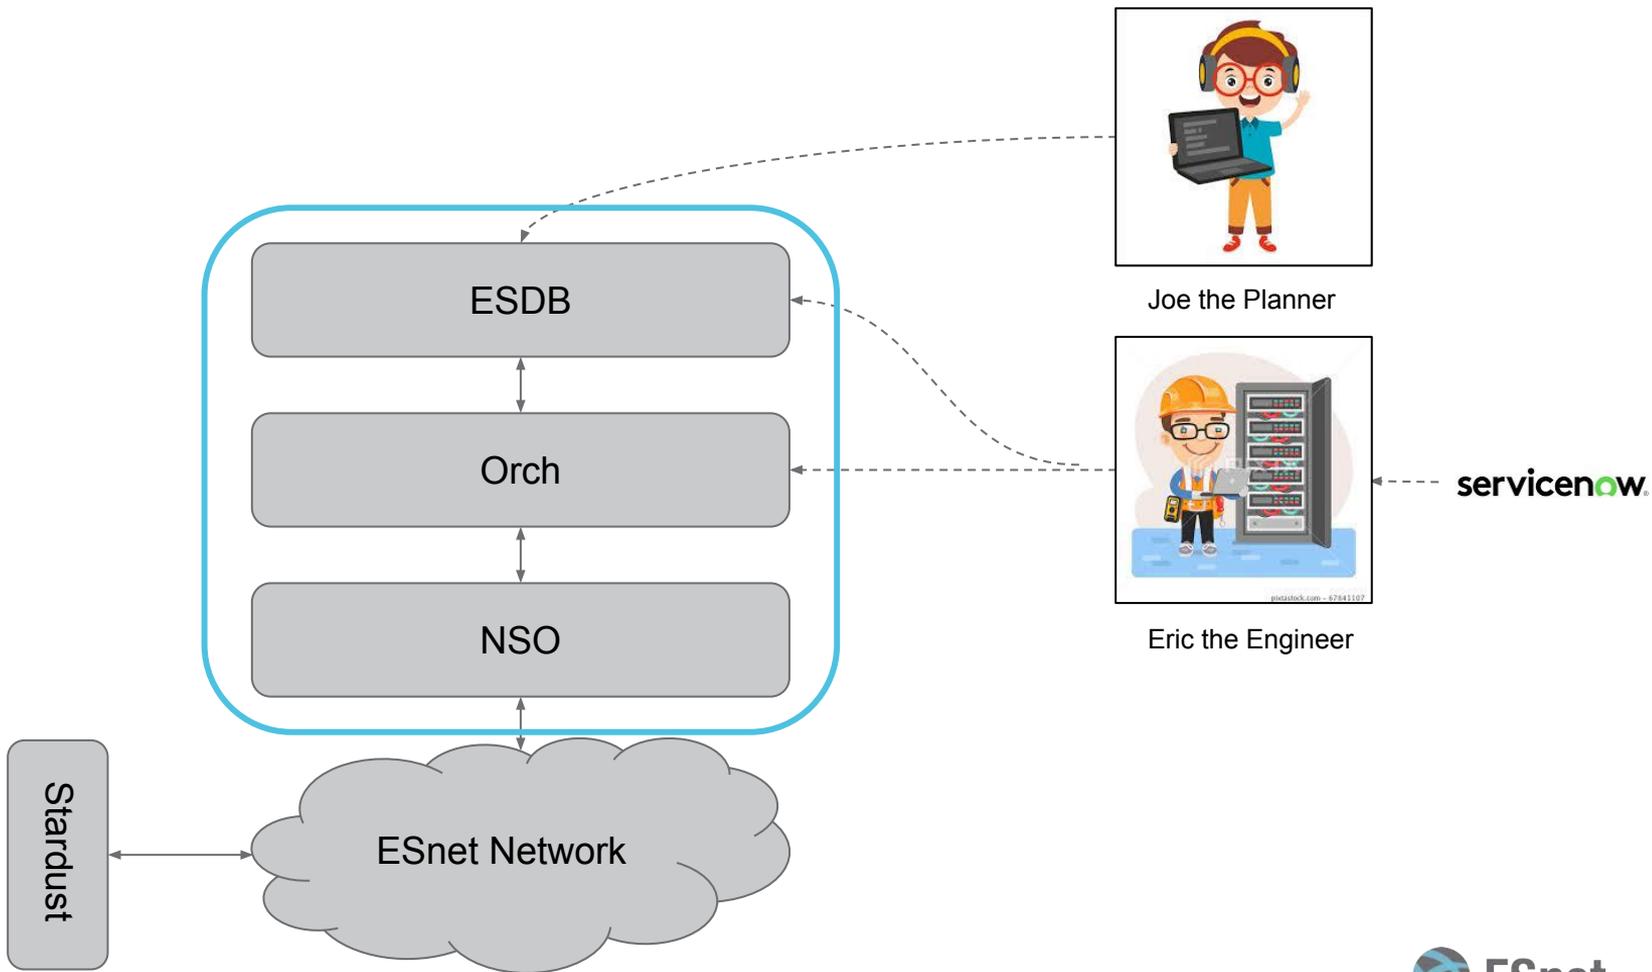
# Why is Orchestration important?

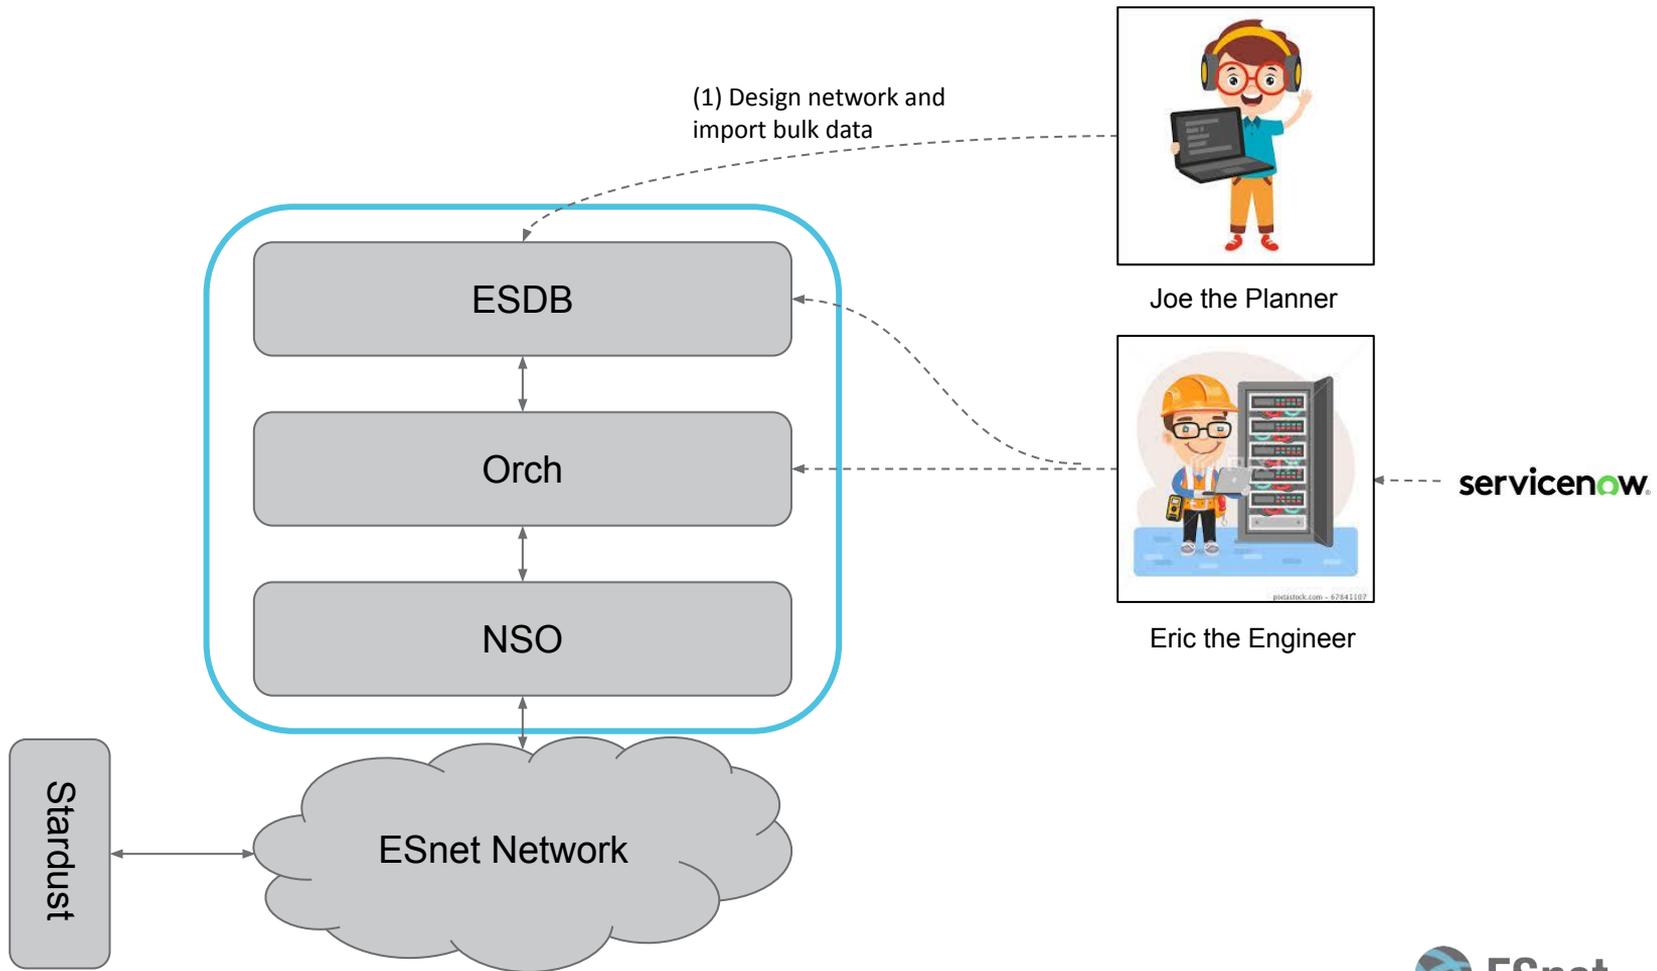
- Consistent and unified configurations
- Eliminate human errors
- Enhance network reliability, resiliency, and robustness
- Allows to manage a larger network
- Heterogeneous equipment
- Better use of human resources
- Planning safety

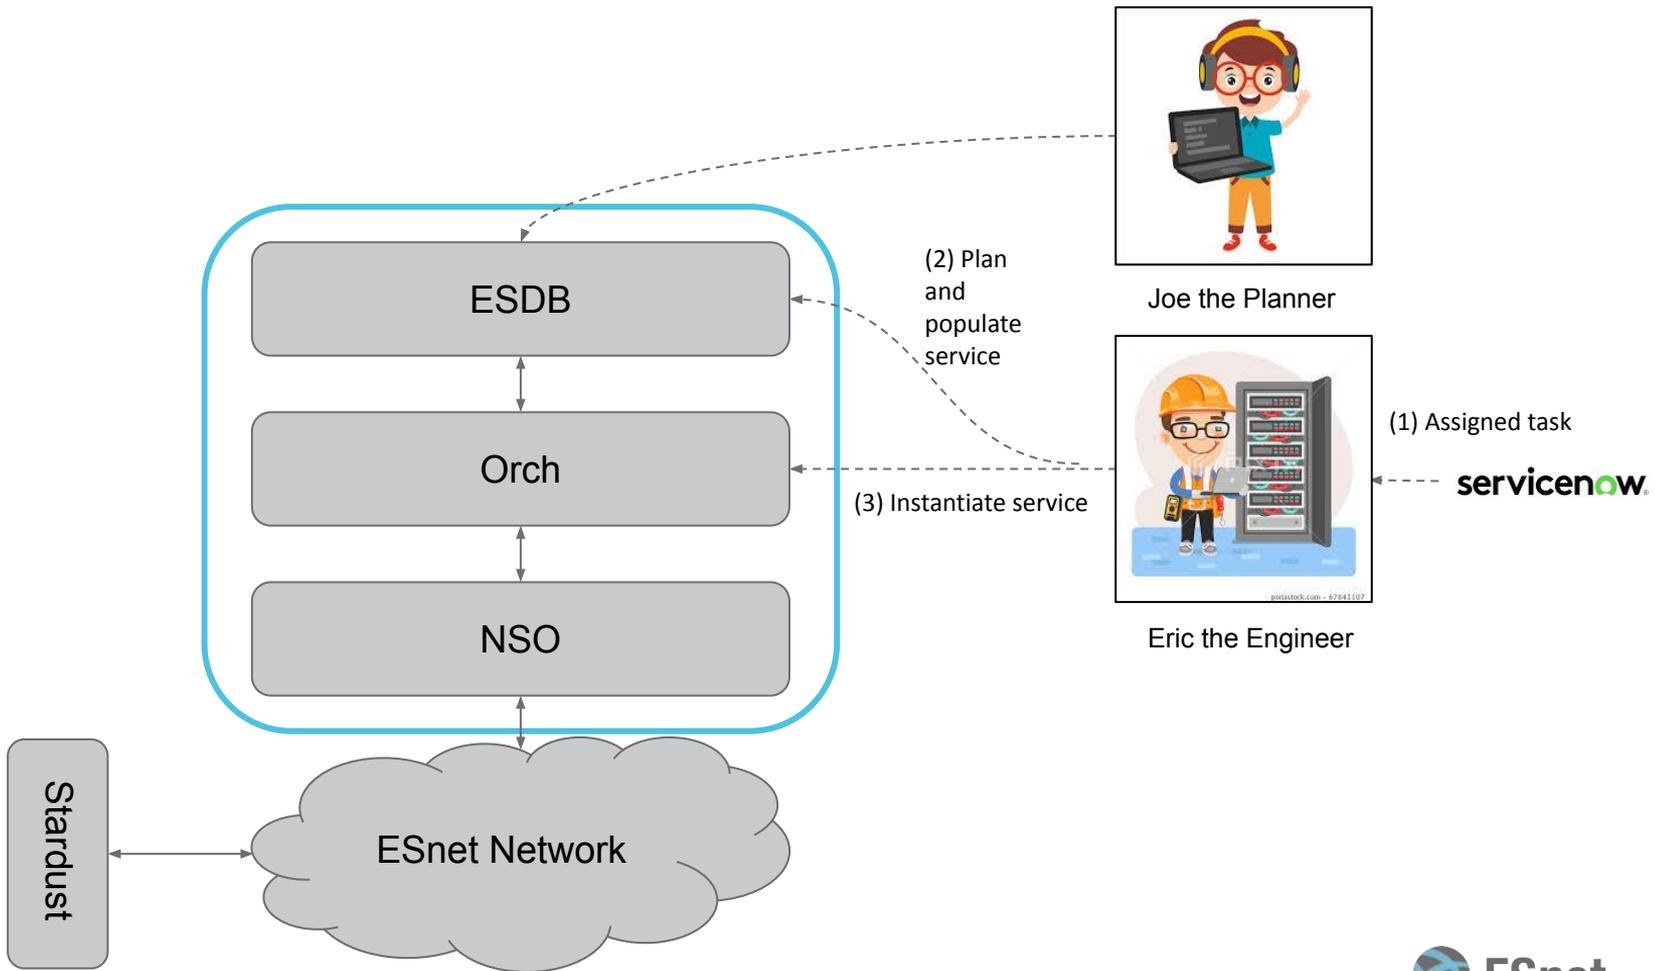
# ESnet Network Orchestration and Automation

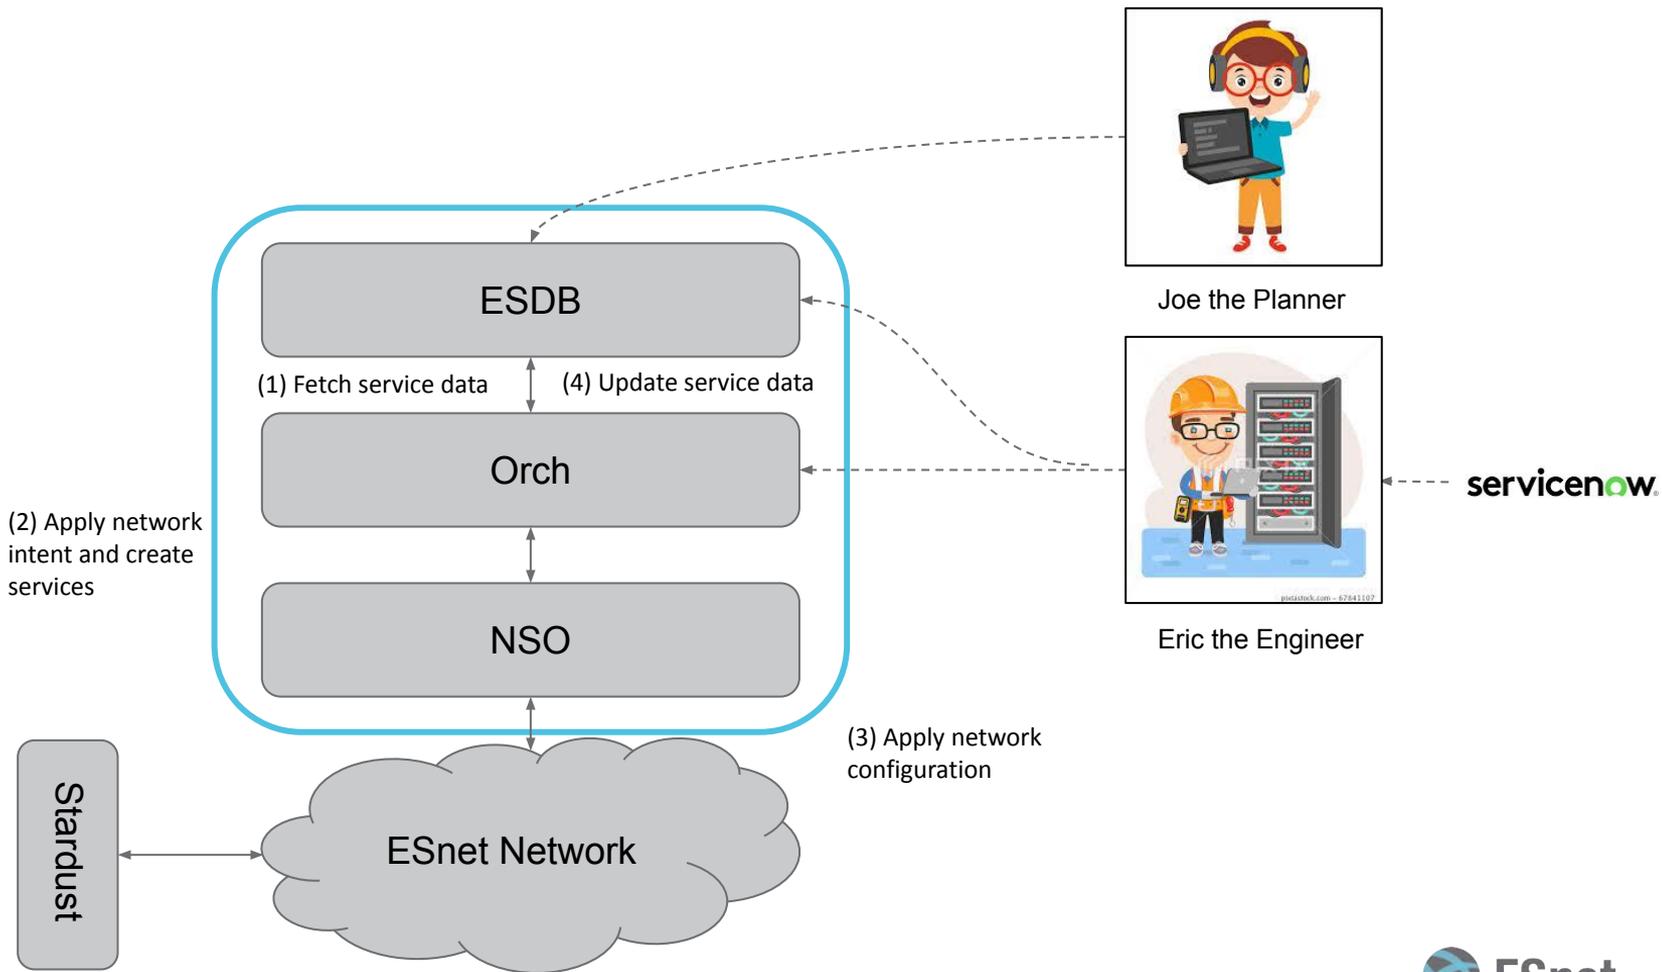


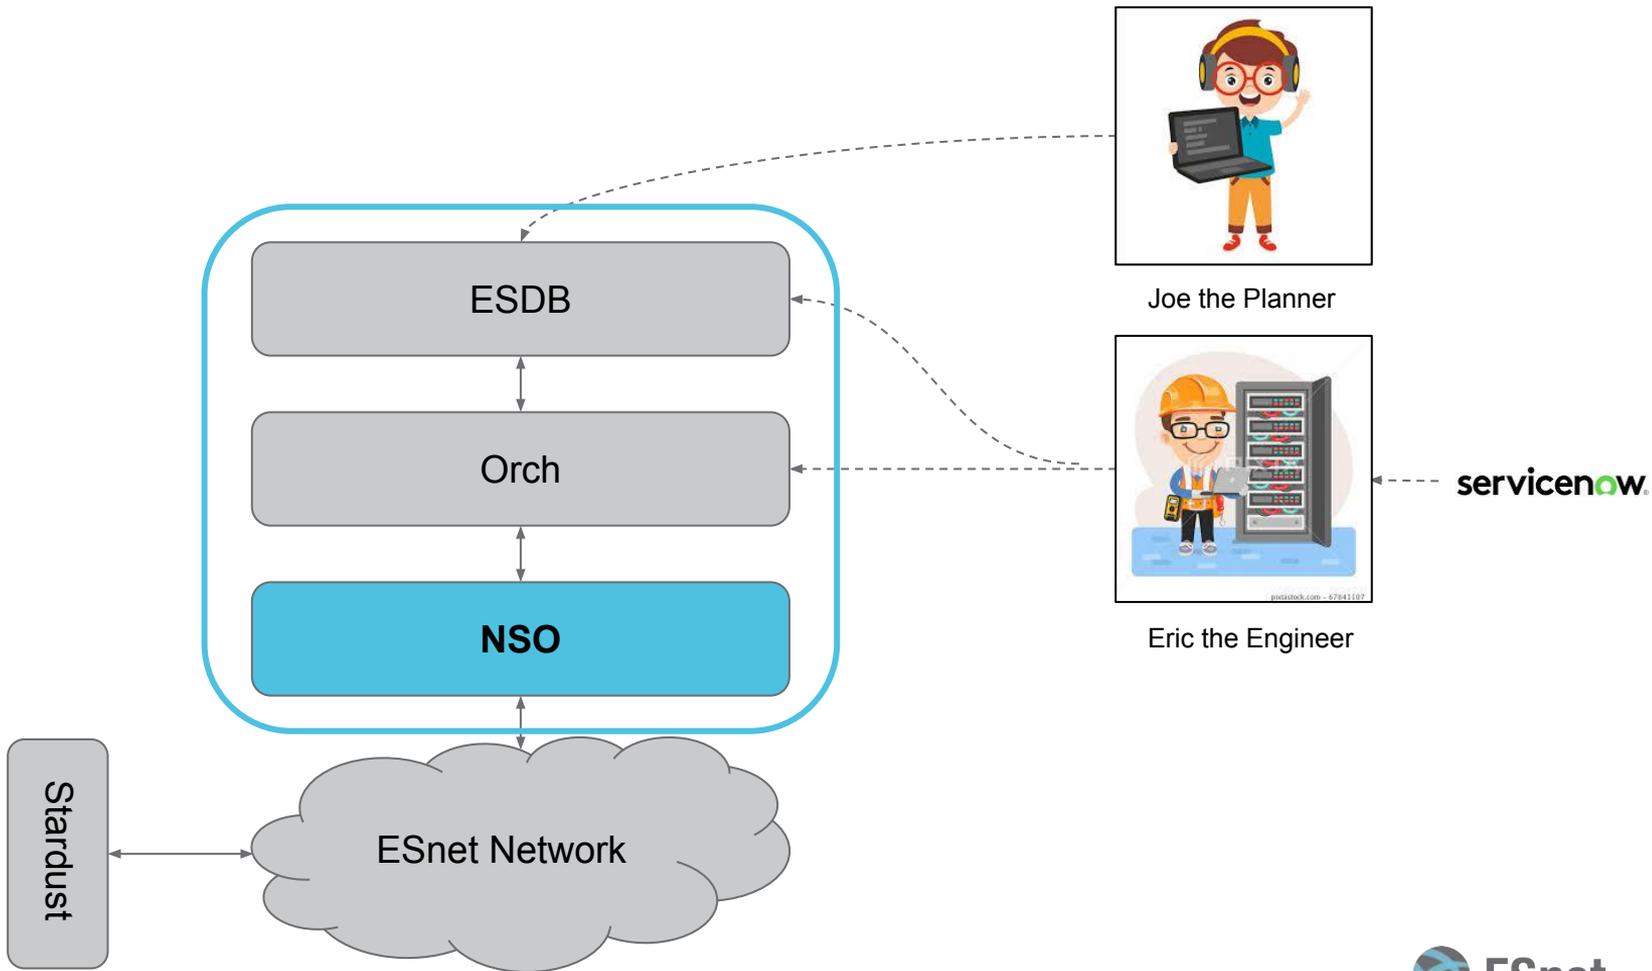


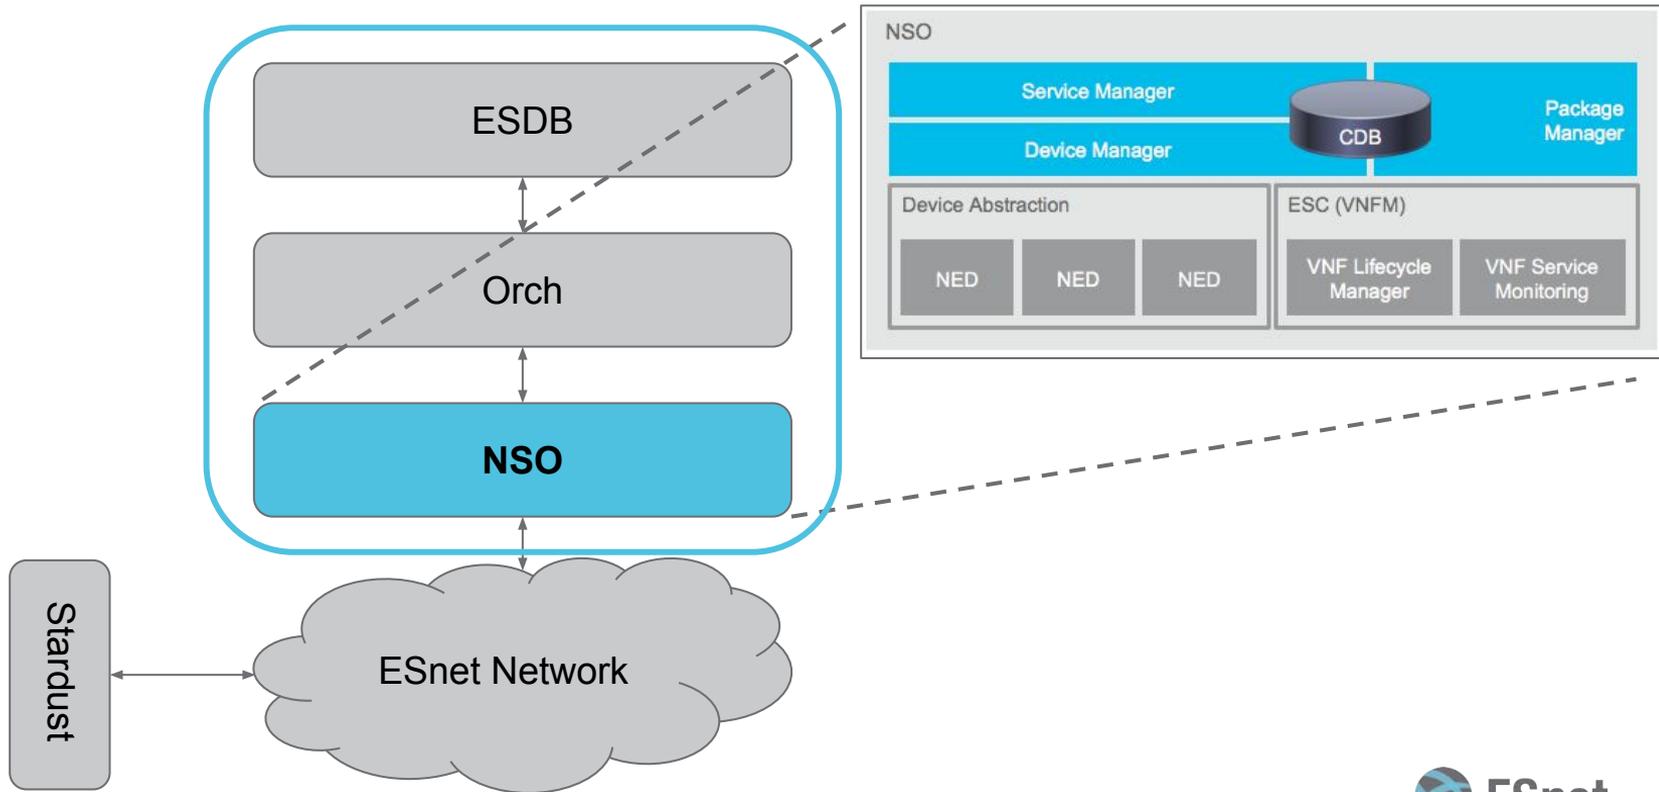








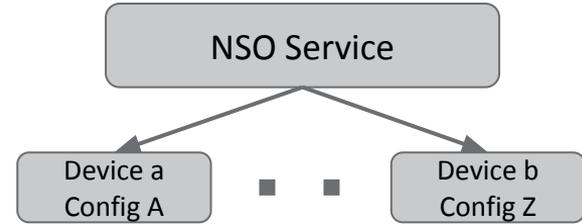




# NSO features used by ESnet

## Architectural features:

- Device agnostic configuration
- Service centric abstraction
  - configure 1, 2, n devices



## Other stuff which comes in handy:

- Data verification
- Service decommissioning and rollbacks

# ESnet's NSO implementation in numbers

Service models: **33**



Services instantiated:

- Devices: 325
- System: 155
- Port: 1705
- BBL: 199
- Bridge: 401
- Host: 161
- LSP: 151
- VPLS: 1
- L3-Interface: 923
- L3-Customer: 277
- L3-Peer: 252
- L3-Transit: 369

# Base-config ask

- Pushing base config to routers
- Combination of vlans, filters, etc.

Personal take on things:

- Pushing base config ahead of time => traditional thinking model
- Intent based networking => we deploy config when we need / use it
- Change of thinking is required IMHO

# Service audit / architecture assessment

- Monolithic services with high complexity
- Validation complexity
  - SAPs are getting validated over multiple services
- Same functionality is getting defined in multiple services
  - IRB interfaces are getting defined in BBL and MPR baseconfig
  - prefix-lists are getting defined in L3, MPR baseconfig, DNS, NTP, syslog, etc.
- Cross dependencies between “higher layer” services
  - Port service has fields for defining its usage - link-type: BBL, VPLS, host, etc.
- Humongous services models and config templates for core services like BBL and L3
- High Test complexity

# Investigate new architectural approach / model with the following objectives

- Clear functional separation and abstraction
- Reduce redundant validation and simplify code complexity
- Dependency hierarchy (dependencies are only pointing downwards)
- Object oriented programming based service approach

# 3-Tier service architecture

1. Services which are directly translate into configuration concepts and build the first device independent abstraction layer
2. Service which build basic service abstraction concepts out the composition of tier 1 services (might introduce additional config as glue between tier 1 services)
3. Services which are composed out of a combination of tier 1 & 2 services and build an operational foundation for ESnet's network operation and customer facing portfolio

# Tier 1 services

Services which are directly translate into configuration concepts and build the first device independent abstraction layer:

- object oriented concept of building blocks
  - leaf-list:devices
- Services must not be able to “function” independently but can
- Examples: Port, prefix-lists, SAPs, SDPs, etc.
- Services expose basically “every” configuration knob
- Namespace: config-services

# General NSO Architecture Revision Concept

Tier 1

interface

vlan

Prefix-lists

SNMP

NTP

netflow

static-routing

DNS

syslog

SSH

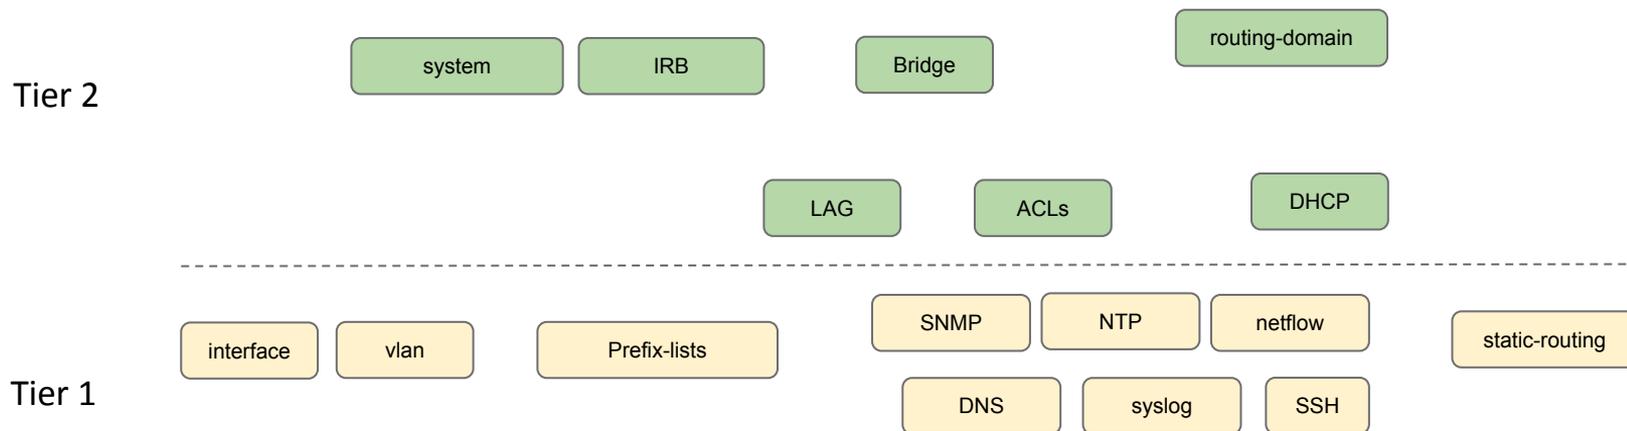


# Tier 2 services

Service which build basic service abstraction concepts out the composition of tier 1 services:

- Composed out of tier-1 services and should build the first layer of functional services which lead to a working service configuration
- Could introduce additional device configuration to glue together tier 1 services and transition them to an independent functional service
- Tier-2 services can be stacked / composed out of other tier-2 services
- Services expose a comprehensive amount of knobs in order to tweak a service configuration
- Namespace: /ncs:base-services

# General NSO Architecture Revision Concept

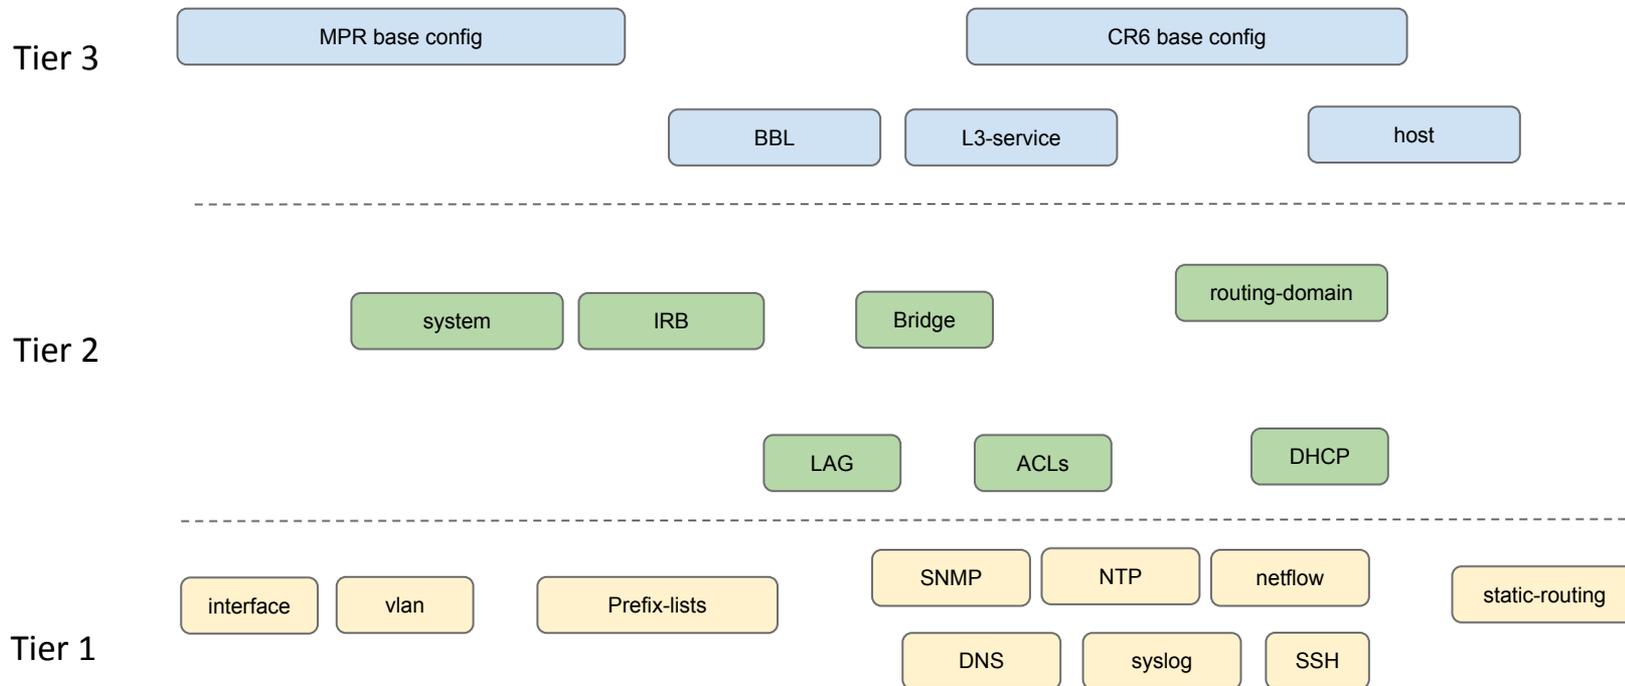


# Tier 3 services

Services which are composed out of a combination of tier 1 & 2 services and build an operational foundation for ESnet's network operation and customer facing portfolio:

- High level service constructs composed out of tier 1 & 2 services
- These services should not have any configuration templating and work only by composing services
- Examples: BBL, L3, MPR & CR6 base config, host, etc.
- Services expose a “minimal” amount of knobs for a rapid service deployment
- Namespace: /ncs:services or /ncs:esnet-services

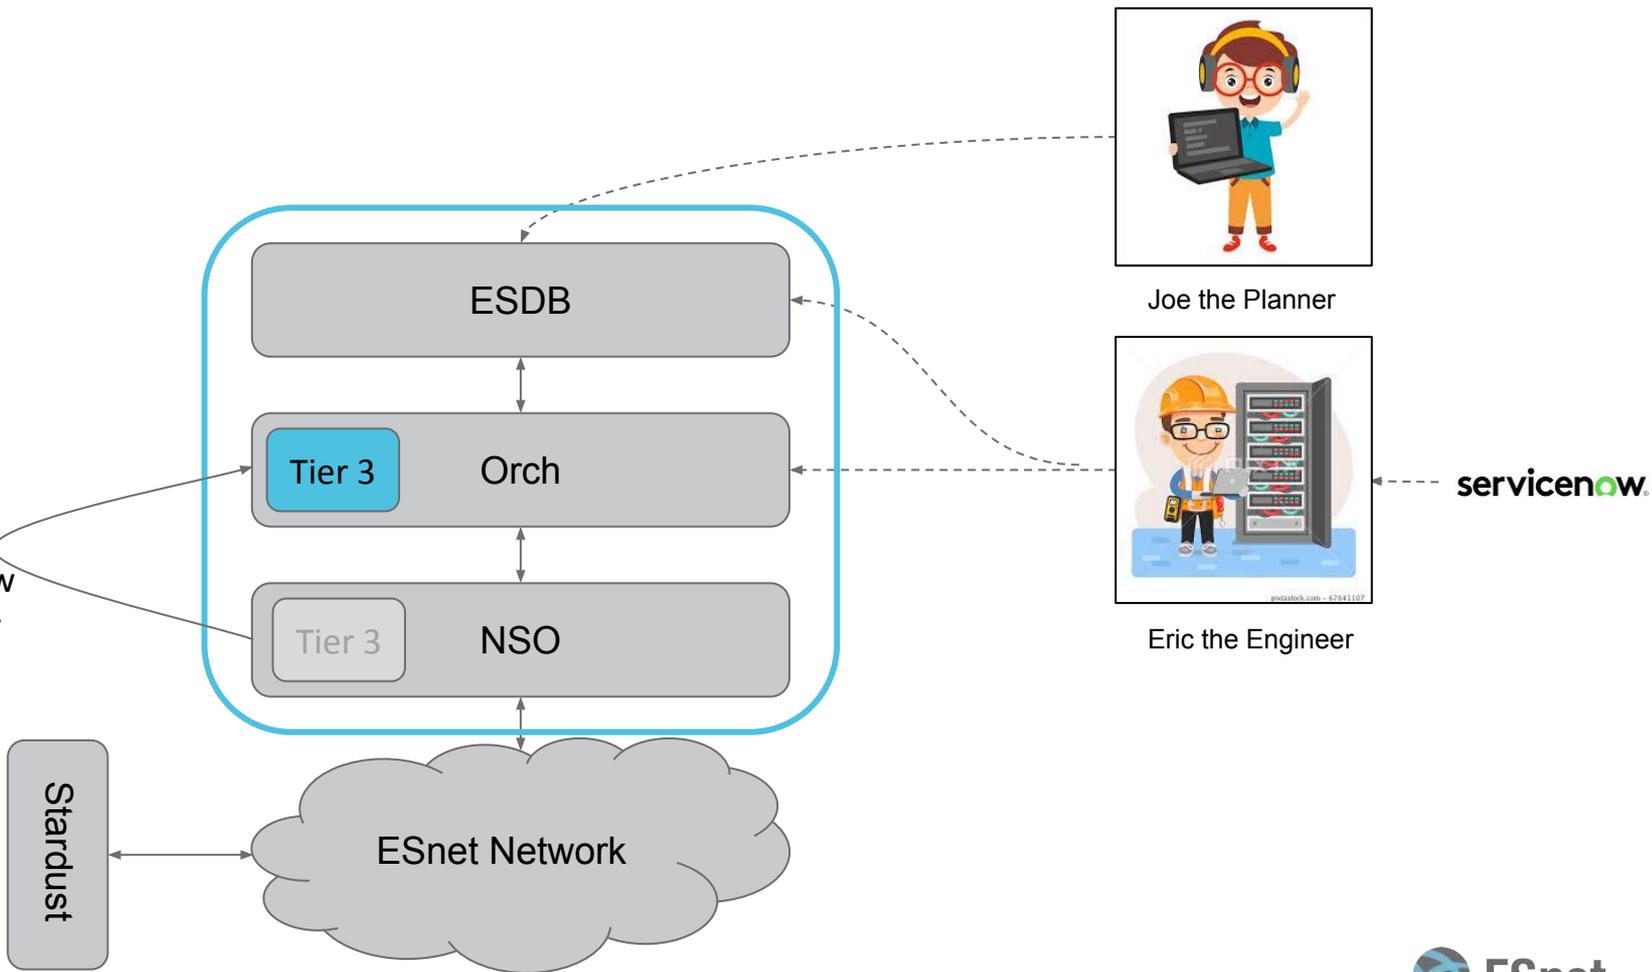
# General NSO Architecture Revision Concept



# Conclusion

- Starting from scratch has too many risks especially since our services are widely used and also intertwined between multiple software layers
- Move tier 3 services into Workflow Orchestrator whenever possible
- Make overall service design a multi layer approach
- Refactor existing services following the tier 2 & 1 service paradigm whenever possible
- Further build out more granular multi-tier service design and architecture

Offloading  
“umbrella”  
services into  
the Workflow  
Orchestrator



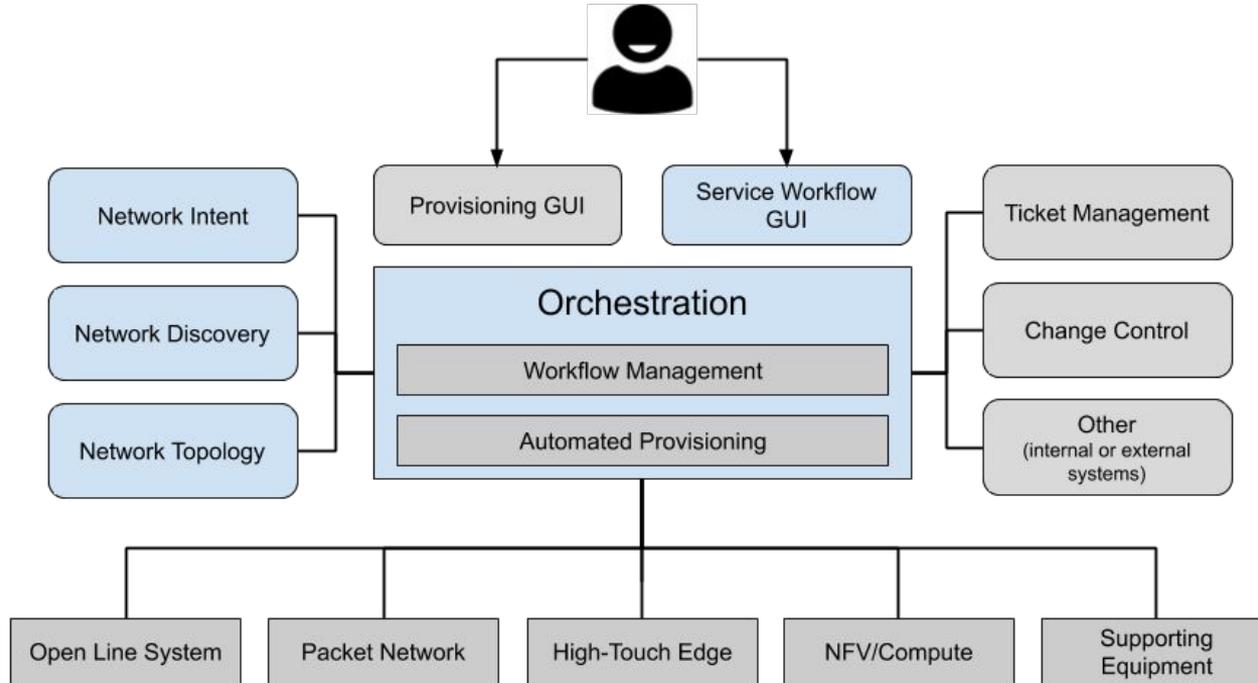
# Future directions

- ESnet is further increasing the amount of orchestrated services
- Use only NSO to manipulate config
  - Reduce inconsistencies
- Transition towards an Workflow Orchestrator only model
  - In a perfect world there won't be any need to touch devices or NSO directly

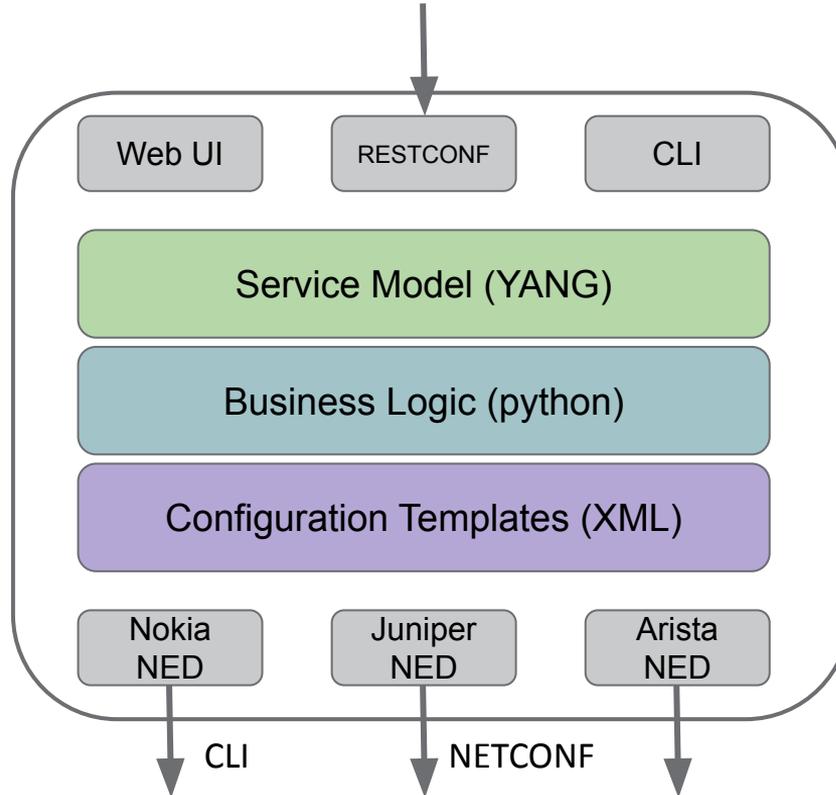
Questions...



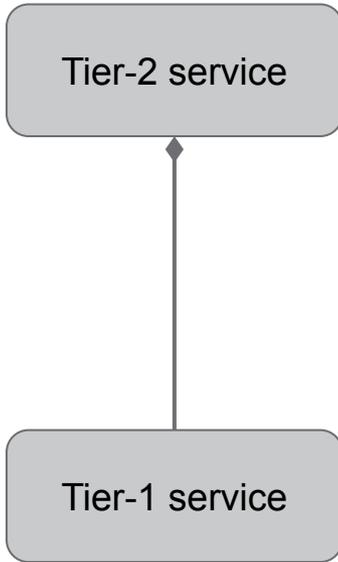
# Orchestration Architecture



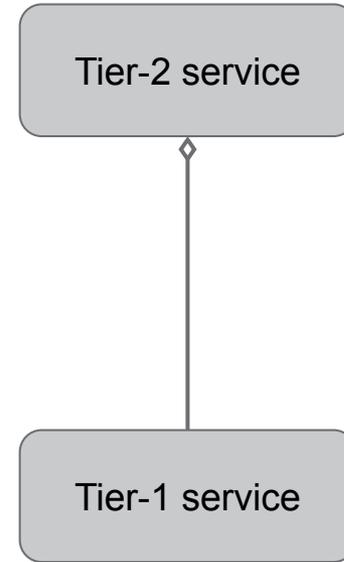
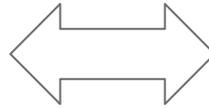
# Anatomy of an NSO service



# Object oriented service architecture design

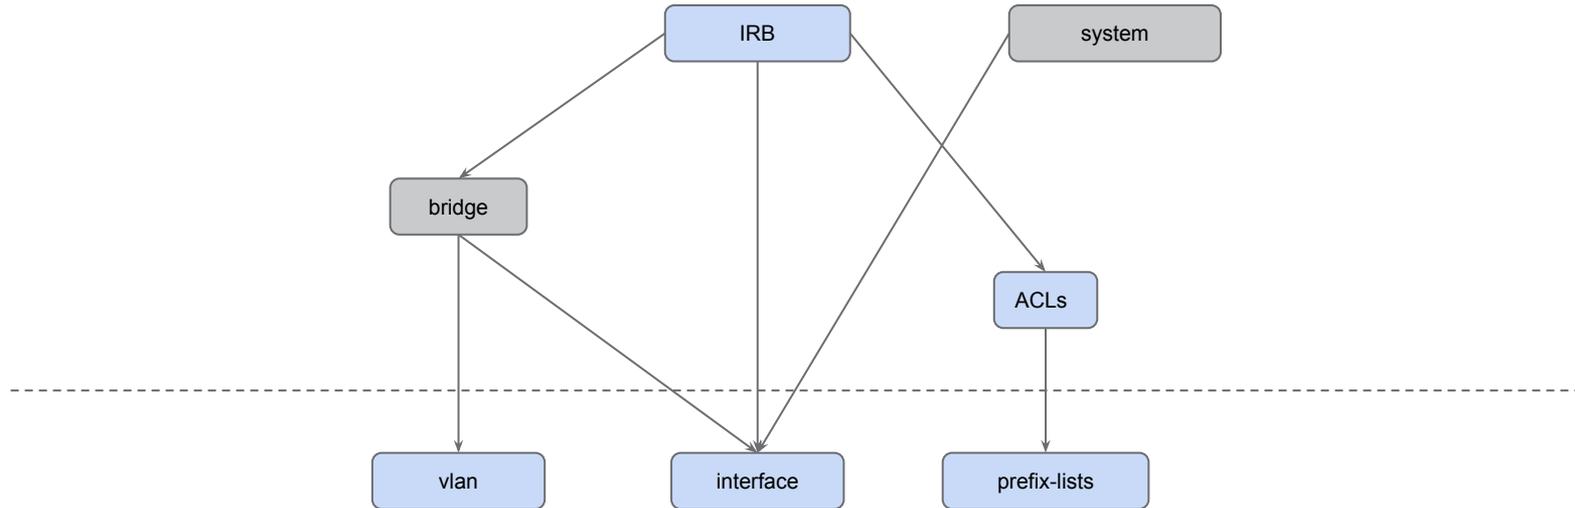


Composition: The tier-2 service creates the tier-1 service instance  
Example:

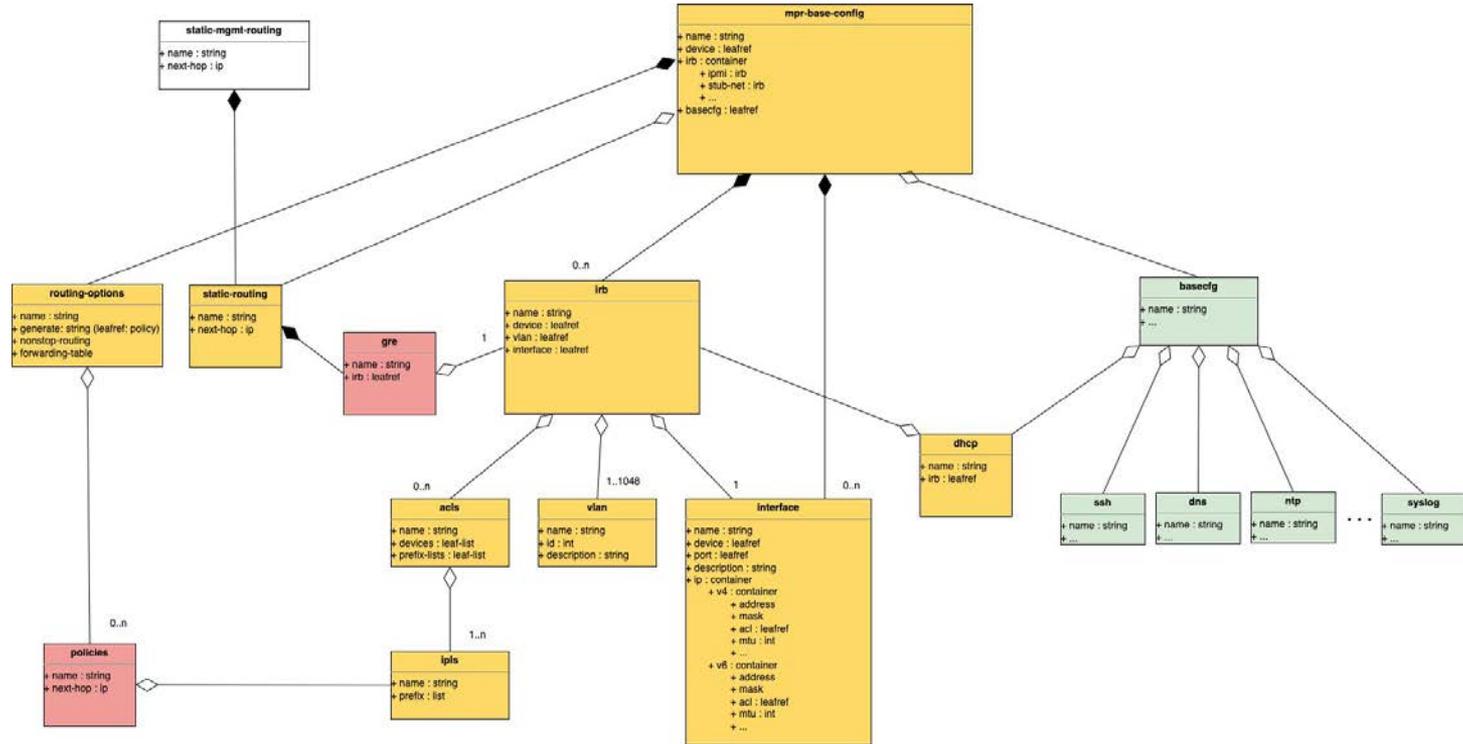


Aggregation: The tier-2 service references the tier-1 service  
Example: multiple IRB (on multiple devices) reference the DNS prefix-list and their ACLs

# Tier-1 to Tier-2 service composition example based on the IRB service



# UML of the Base-configuration services



# ESDB GUI

 **ESnet**  
ENERGY SCIENCES NETWORK

**ESDB** **ESnet**

HOME ›

## Circuits

<b>ESnet Circuit ID</b>	<b>Endpoint A/Z</b>	<b>Orchestrator ID</b>	<b>Provider / ID</b>
<a href="#">ESNET-Z00227</a> In-service	torrestal-mpr1:ge-0/0/40 DOE-NNSA-FORR:Site_conn		<a href="#">DOE-IN Headquarters Forrestal Building</a>
<a href="#">ESNET-Z00330</a> In-service	ATLA-Ciena-Shelf2:ETH10G-2-7-1 ATLM-Ciena-Shelf1:ETH10G-1-7-1		<a href="#">Internet2</a> ESNET-ATLA-ATLM-10GE-22437
<a href="#">ESNET-Z00349</a> In-service	ATLA-Ciena-Shelf2:ETH10G-2-7-4 ATLM-Ciena-Shelf1:ETH10G-1-7-4		<a href="#">Internet2</a> ESNET-ATLA-ATLM-10GE-22440
<a href="#">ESNET-Z00355</a> In-service	ATLA-Ciena-Shelf2:ETH10G-2-7-3 ATLM-Ciena-Shelf1:ETH10G-1-7-3		<a href="#">Internet2</a> ESNET-ATLA-ATLM-10GE-22439

- +** ADD CIRCUIT
- 🔧** RESERVE PORT
- 📄** DOWNLOAD CIRCUIT LIST

**Circuits**

Couplers

# Orch GUI

The screenshot displays the Orch GUI interface. At the top, the Orchestrator logo is on the left, and the status 'Engine is Running' with a green dot, a red checkmark, and a notification icon with '4' is on the right. Below the navigation bar (Processes, Subscriptions, Metadata, Tasks, Settings), there is a 'New Process' button. The main content is divided into two sections: 'Active Processes' and 'Completed Processes'. Both sections have a search bar and navigation arrows. The 'Active Processes' table shows three failed system processes. The 'Completed Processes' table shows four successful processes.

Assignee	Status	Workflow	Target	Abbr.	Subscription(s)	Started	Modified
SYSTEM	Failed	create_prefix_list	CREATE	ESNET	AS3443 - ESNET - Energy Sciences Network -- ipt	10-7-2023	24-8-2023
SYSTEM	Failed	create_prefix_list	CREATE	ESNET	AS65152 - ESNET - Energy Sciences Network -- ipt	2-6-2023	9-6-2023
SYSTEM	Failed	create_l3_service	CREATE	ESNET	L3-transit for 65512--base--lbnl59qa-cr6a:esnet_lo-685	8-5-2023	2-6-2023

Workflow	Target	Customer	Product(s)	Tag(s)	Subscription(s)	Modified
create_physical_connection	CREATE	FABRIC:FABRIC Testbed	Physical Connection	PCS	FABRIC conn-ak: star-mpr1:ge-0/0/42 1G	12:18 CET
provision_node_enrollment_tpd	MODIFY	ESNET:Energy Sciences Network	Enroll Transponder	NESTPDR	Node eqxda3-tpdr5 Provisioned	08:43 CET
provision_node_enrollment_tpd	MODIFY	ESNET:Energy Sciences Network	Enroll Transponder	NESTPDR	Node seat-tpdr4 Provisioned	08:35 CET
modify_backbone_link_member_state	MODIFY	ESNET:Energy Sciences Network	ECMP Group	BBL	ga-cr6 to sunn-cr6	6-9-2023

