Status telemetria e monitoraggio

Giovanni Cesaroni, Nino Ciurleo

GARR

Agenda

- **GARR-T monitoring**: API layer del BSS
- **GARR-T logging**: low code log exporter
- **GARR-T flows**: nuovo sistema OLAP

GARR-T monitoring: dove eravamo

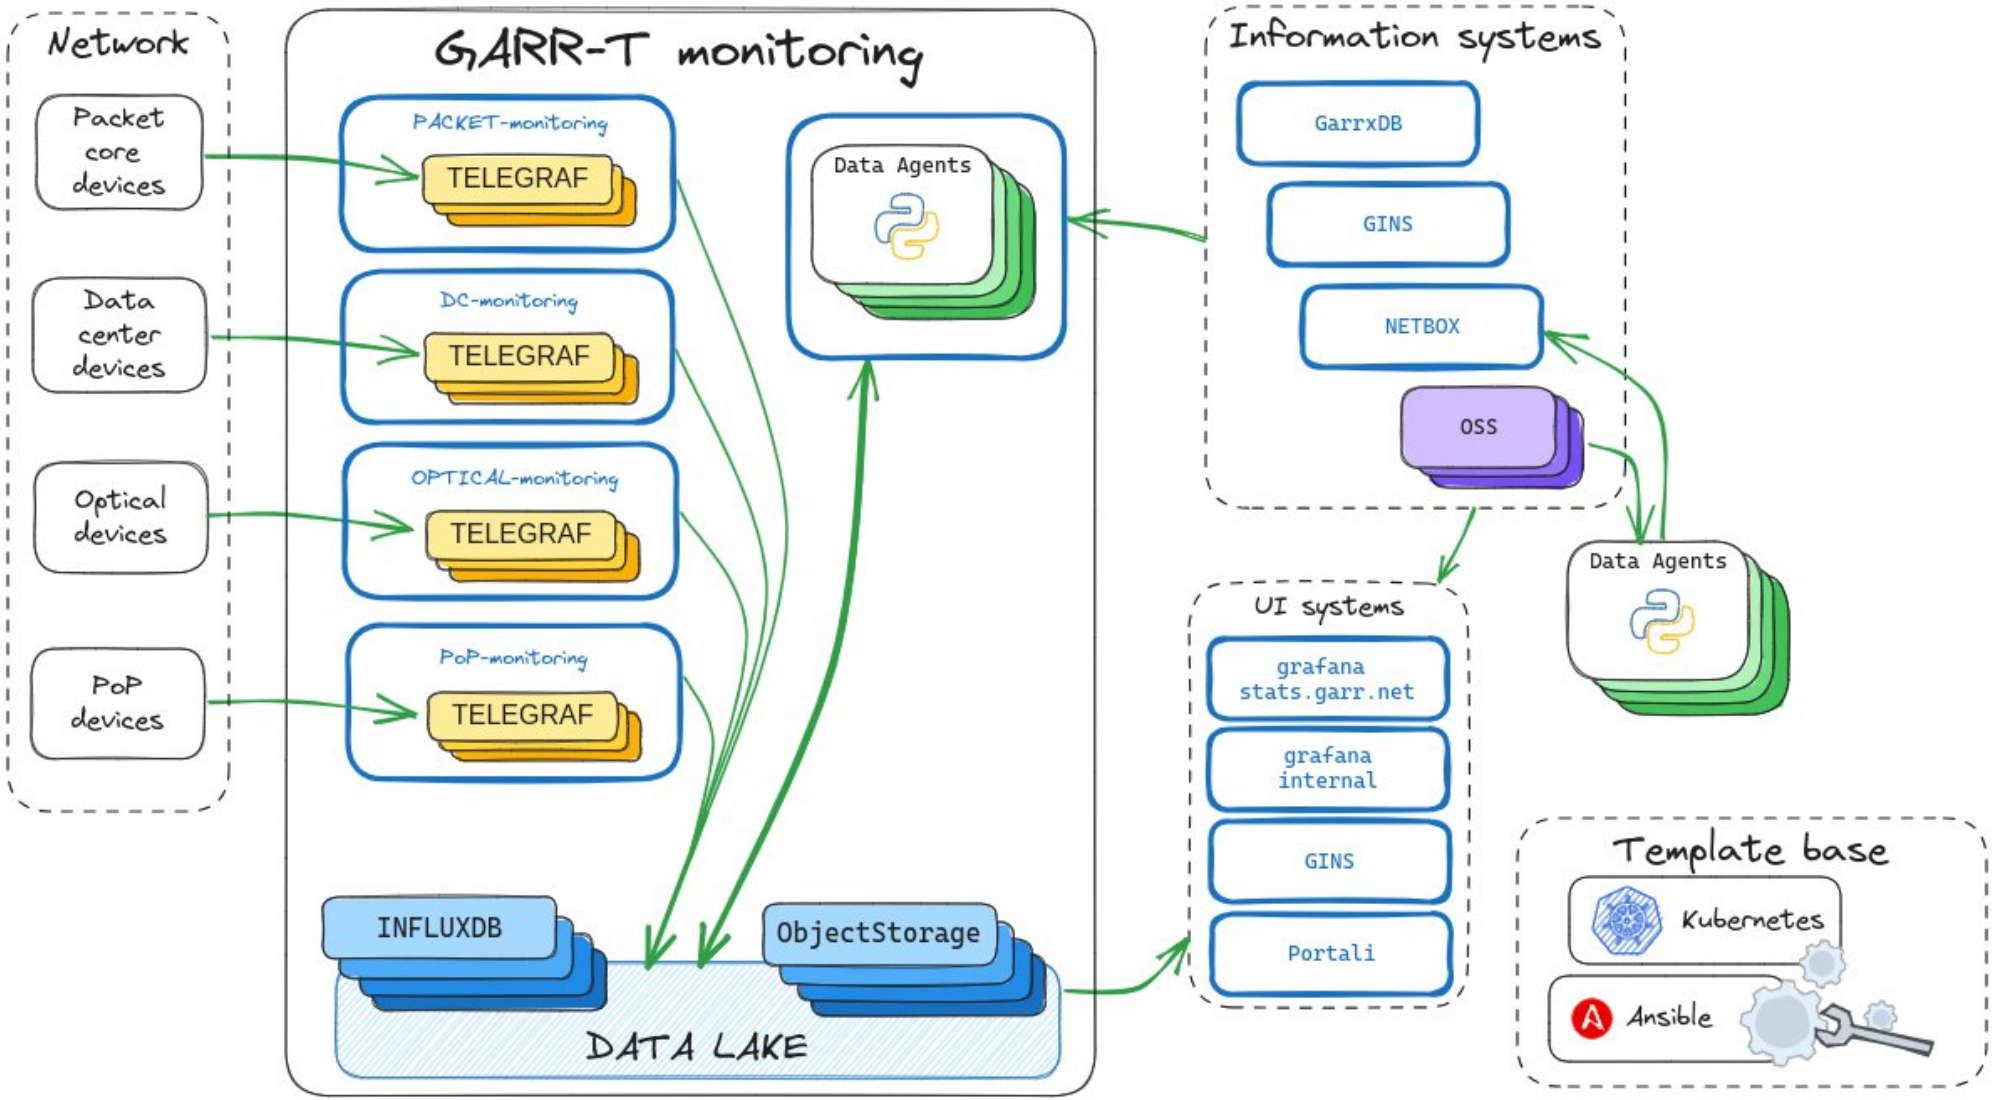

GARR-T monitoring: dove siamo

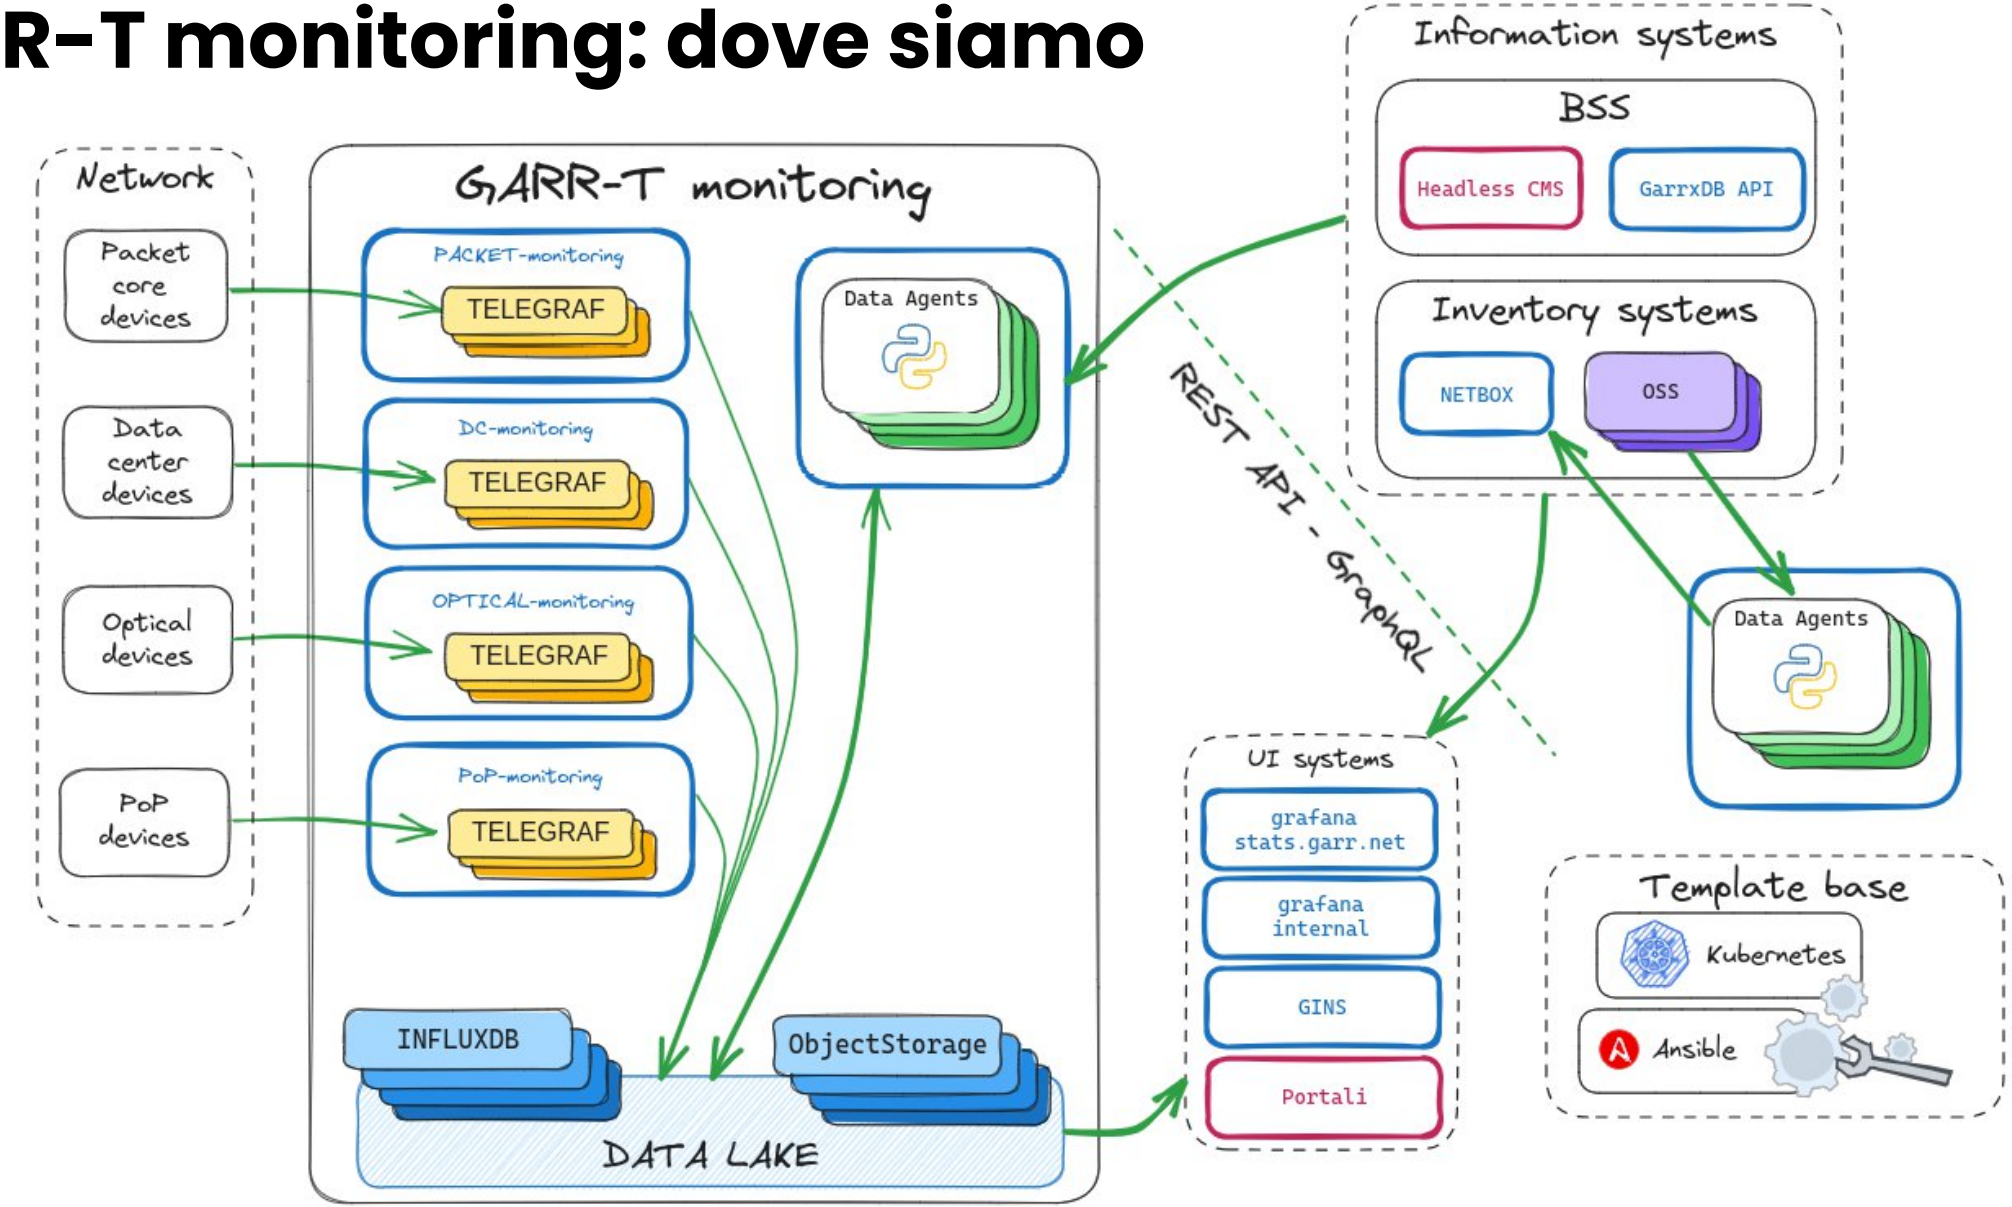

Requisiti API layer per BSS

1. **Automazione**: **API REST**, OpenApi, GraphQL, Swagger
2. **Autorizzazione**: **RBAC** per utenti + API token per applicazioni
3. **Autenticazione**: integazione con **SSO**
4. **Velocita'** nell'implementazione di **nuovi modelli**
5. **Open source**
6. **Documentazione**

Strapi: cos'e' e come lo usiamo?

1) Headless CMS designed for **customization**
2) Soddisfa i requisiti

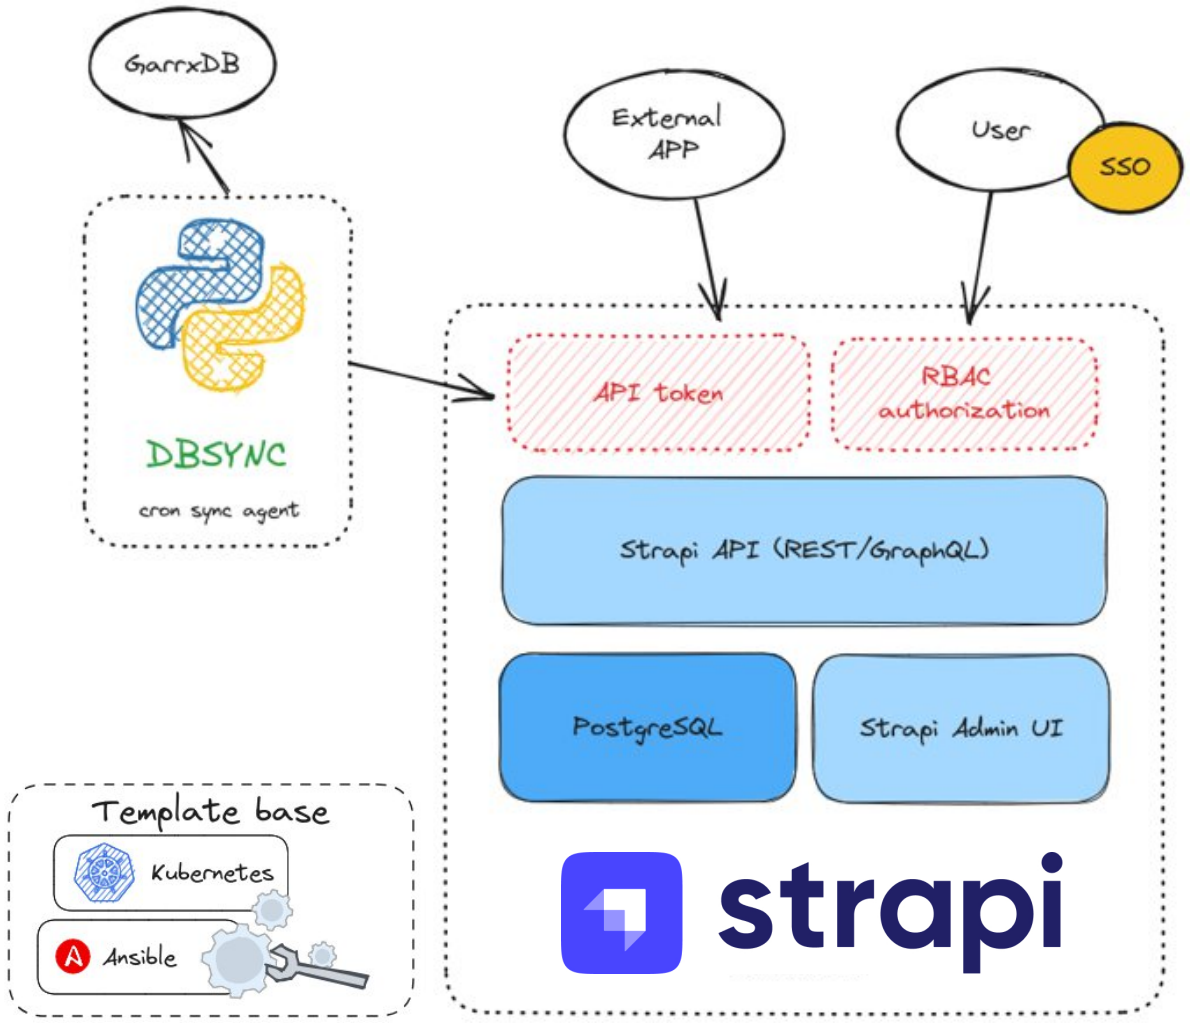

primo caso d'uso: classificazione degli enti

Institutions, network service count

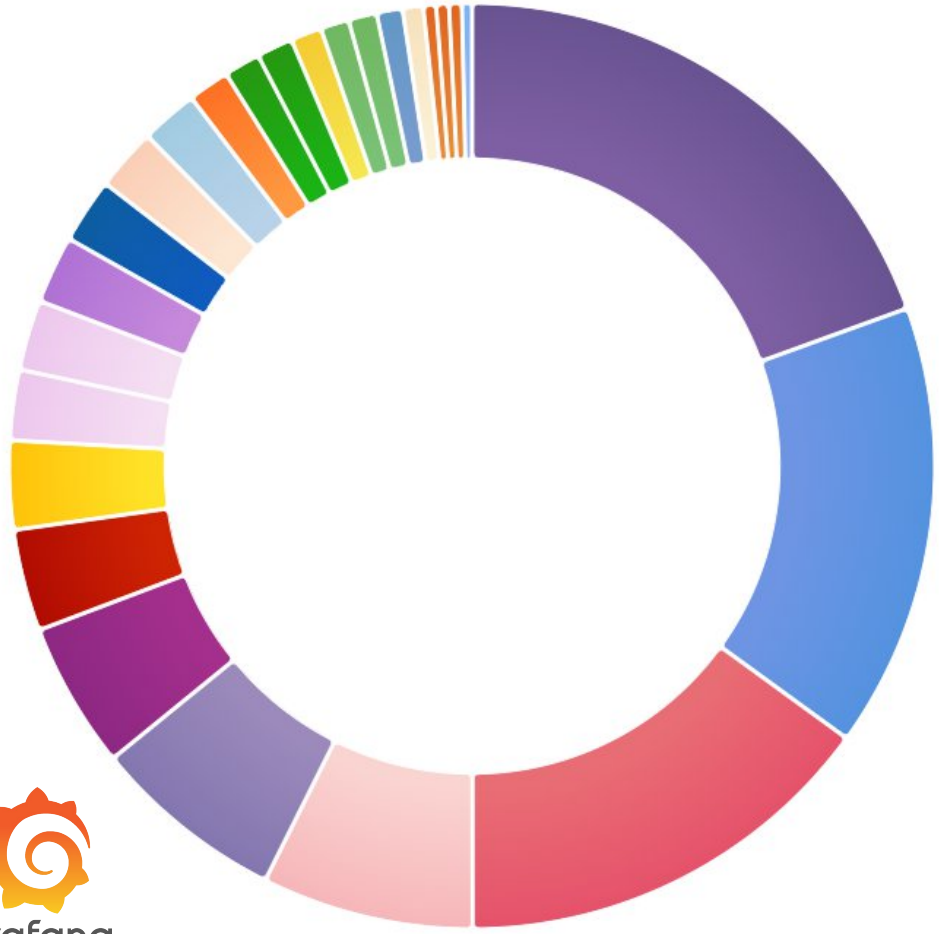

	Value
CNR	212
SCUOLE SECONDARIE	169
UNIVERSITÀ	163
INFN	80
ENTI DI RICERCA SCIENTIFICA E TECNOLOGICA	73
IRCCS	55
ISTITUTI INTERNAZIONALI	38
CONSERVATORI	33
UNIVERSITÀ NON STATALI	26
INAF	26
ARCHIVI	25
BIBLIOTECHE	24
ENEA	23
INGV	21
ACCADEMIE	15
ASI	13
ENTI DI FORMAZIONE	13
CONSORZI UNIVERSITARI	11
MUSEI E ISTITUTI CULTURALI	10
IZS	10
ISTITUTI RICERCA BIOMEDICA	9
AMMINISTRAZIONE PUBBLICA	7
SCUOLE PRIMARIE E MEDIE	4
ENTI DI PROMOZIONE DELLA SCIENZA	4
MUSEI SCIENTIFICI	4
ISTITUTI D'ARTE	3

Grafana

Access type count

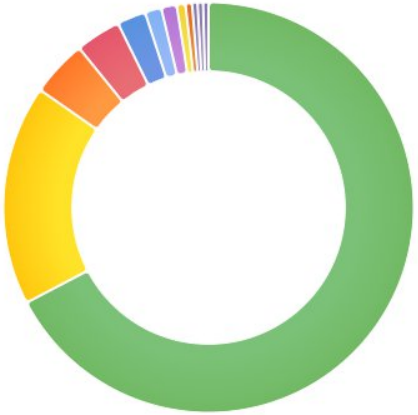

	Value
DF	770
L3	198
L2VPN EoMPLS	50
LAMBDA	40
L2VPN	25
L2VPN MPLS	14
L2	13
PonteRadio	7
OPN	5
CDN	4
L3VPN MPLS	4
L3VPN	4

Service type count

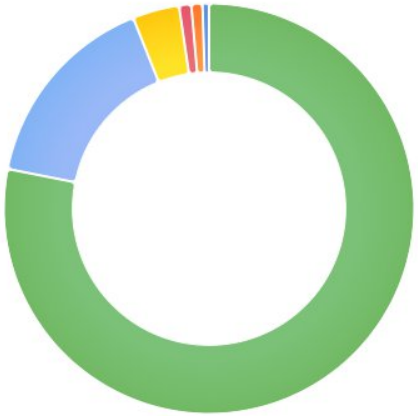

	Value
IP	889
VPN	180
PEERING	41
OPN	10
GARR	9
AGGREGAZIONE	5

GARR-T logging

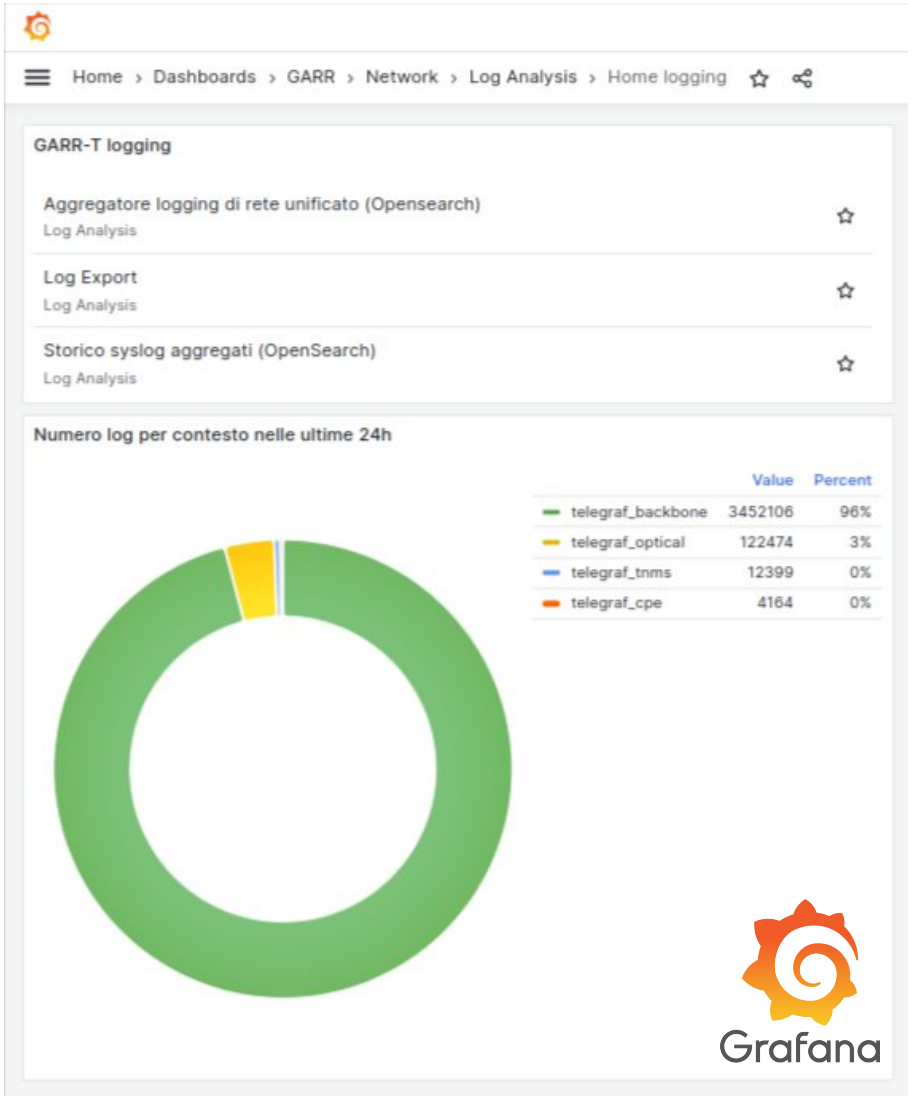

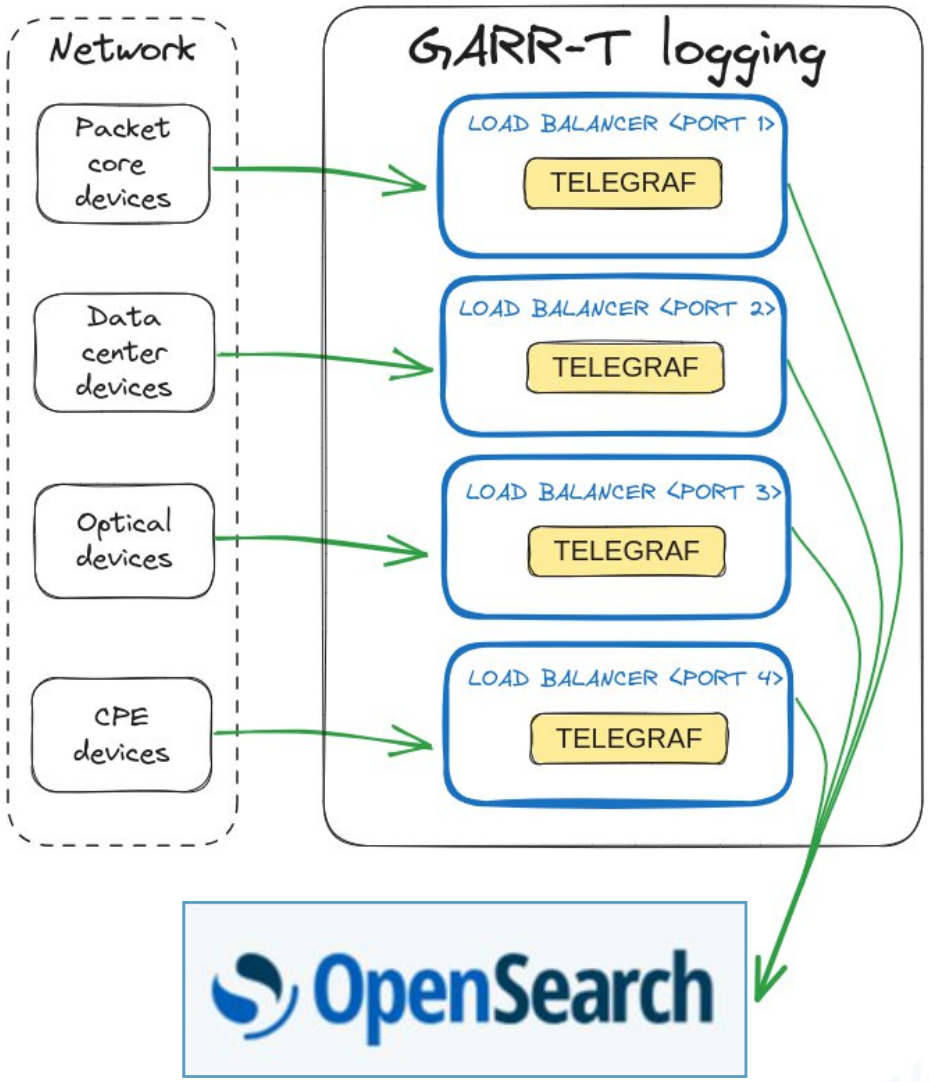

GARR-T logging: low code log exporter

GARR-T logging: low code log exporter

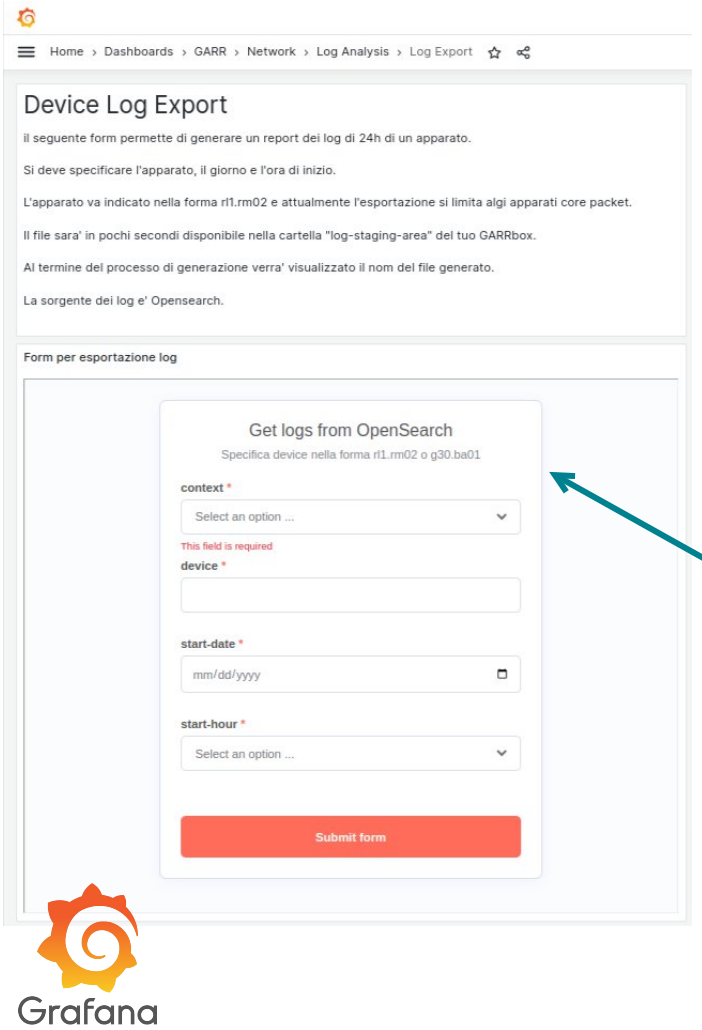

Grafana

OpenSearch

n8n

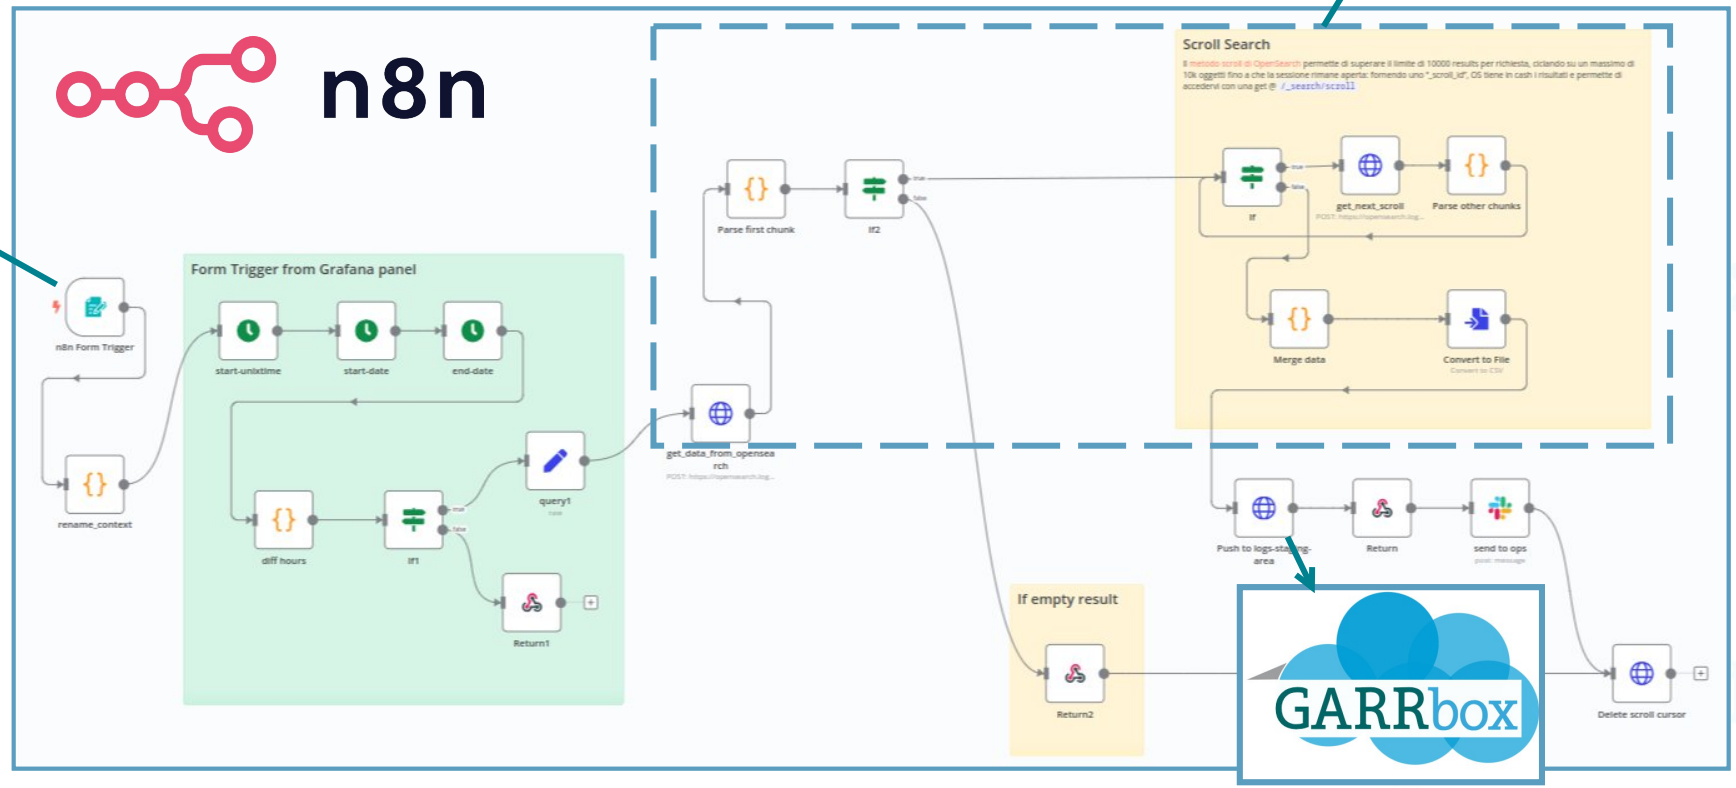

WORKSHOP GARR 2024

NET MAKERS

Monitoraggio dei flussi

Monitoraggio dei flussi: agenda

- Architettura/implementazione sistema OLAP Netflow
- Data flow
- Aggregazione dati
- Casi d’uso
- Conclusioni

Architettura/implementazione

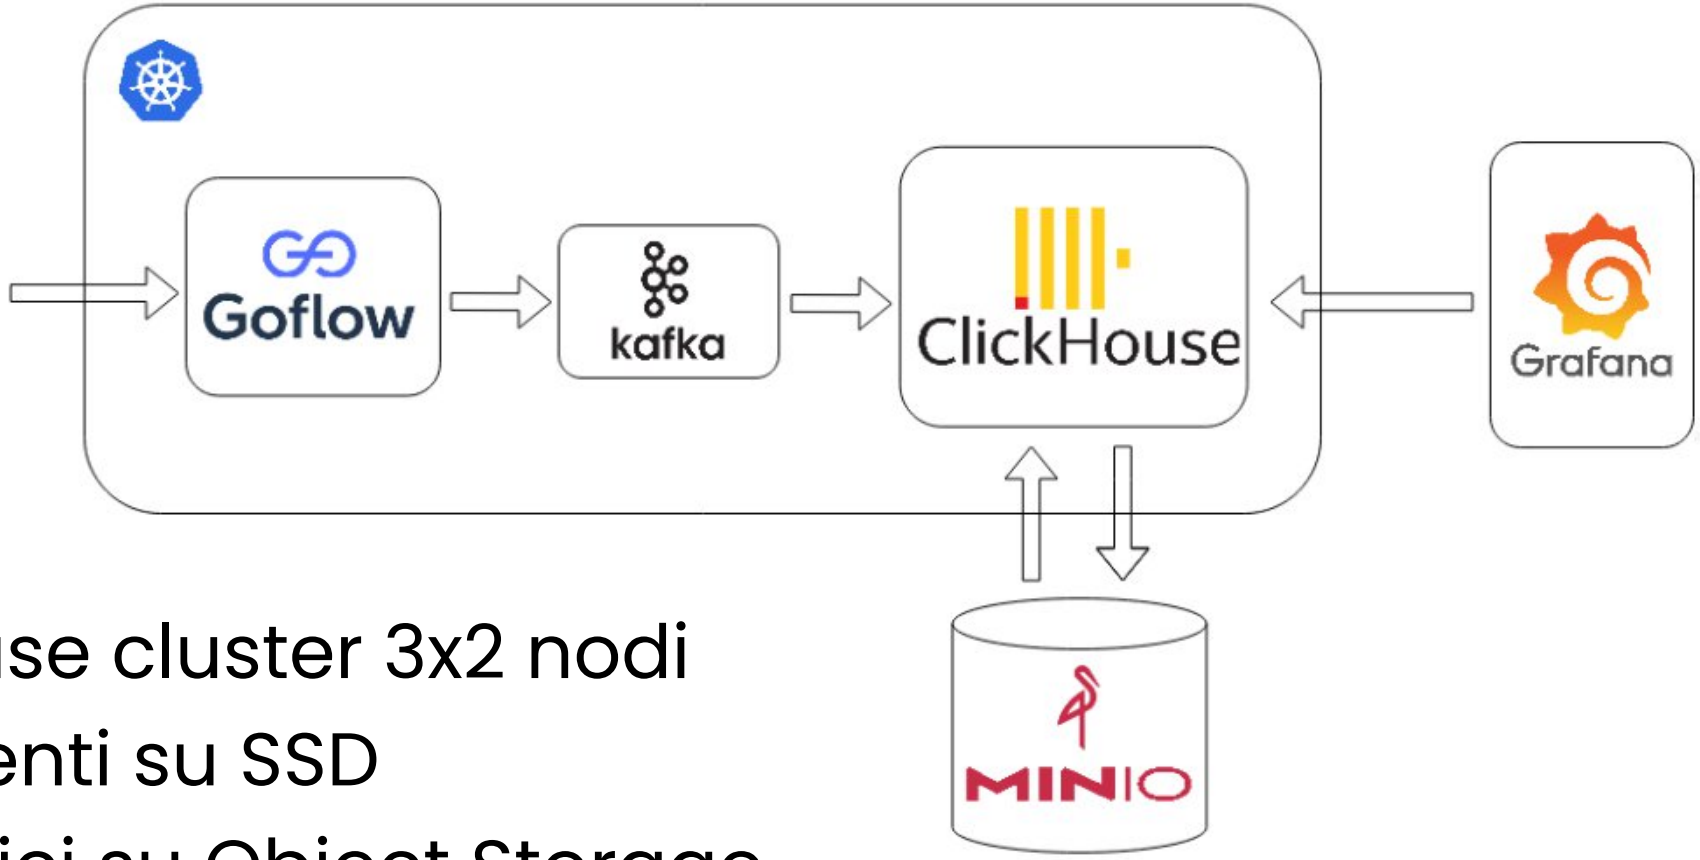

- Clickhouse cluster 3x2 nodi
- Dati recenti su SSD
- Dati storici su Object Storage

Data flow

Netflow

`TimeReceived` UInt64,
`TimeFlowStart` UInt64,
`TimeFlowEnd` UInt64,
`SequenceNum` UInt32,
`InIf` UInt32,
`OutIf` UInt32,
`SrcAS` UInt32,
`DstAS` UInt32,
`EType` UInt32,
`Proto` UInt32,
`SrcPort` UInt32,
`DstPort` UInt32,
`Bytes` UInt64,
`Packets` UInt64,
`SamplerAddress` FixedString(16),
`SrcAddr` FixedString(16),
`DstAddr` FixedString(16),
`NextHop` FixedString(16),
`IPTos` UInt32,
`IPTTL` UInt32,
`TCPFlags` UInt32,
`IcmpType` UInt32,
`IcmpCode` UInt32

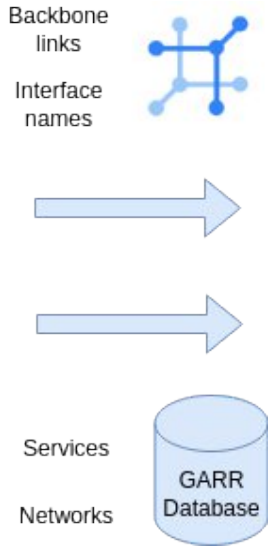

Decorati

`TimeReceived` DateTime,
`TimeFlowStart` DateTime,
`TimeFlowEnd` DateTime,
`InIf` UInt32,
`OutIf` UInt32,
`InIfName` LowCardinality(String),
`OutIfName` LowCardinality(String),
`SrcAS` UInt32,
`DstAS` UInt32,
`EType` UInt32,
`SrcPort` UInt32,
`DstPort` UInt32,
`Bytes` UInt64,
`Packets` UInt64,
`v6SamplerAddress` IPv6,
`sampler` LowCardinality(String),
`v6SrcAddr` IPv6,
`source` LowCardinality(String),
`src_service` LowCardinality(String),
`src_network` IPv6,
`src_network_mask` UInt8,
`v6DstAddr` IPv6,
`destination` LowCardinality(String),
`dst_service` LowCardinality(String),
`dst_network` IPv6,
`dst_network_mask` UInt8,
`in_isis` Bool,
`out_isis` Bool,
`in_lsi` Bool,
`out_lsi` Bool,
`v6NextHop` IPv6,
`IPTTL` UInt32,
`Proto` UInt8,
`IPTos` UInt8,
`TCPFlags` UInt8,
`IcmpType` UInt8,
`IcmpCode` UInt8,
`partition` UInt8,
`is_subnet_src` Bool,
`is_subnet_dst` Bool

Filtrati ed Aggregati

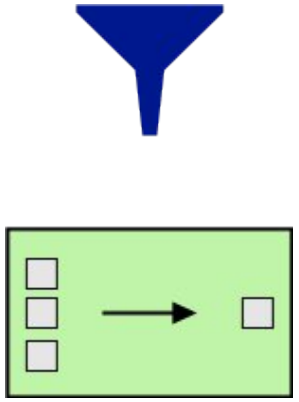

`time` DateTime,
`sampler` LowCardinality(String),
`is_subnet_src` Bool,
`is_subnet_dst` Bool,
`source` LowCardinality(String),
`destination` LowCardinality(String),
`src_serviceID` UInt16,
`dst_serviceID` UInt16,
`InIf` UInt32,
`OutIf` UInt32,
`SrcAS` UInt32,
`DstAS` UInt32,
`Packets` UInt64,
`Bytes` UInt64,
`Flows` UInt64

Aggregazione dei dati

Clickhouse ha un sistema di aggregazione per ottenere tabelle con dati **filtrati**

Tabella **user_id**:

- Matrice di traffico end-to-end tra Sedi/AS
- Esclusione flussi di backbone
- Aggregazione giornaliera

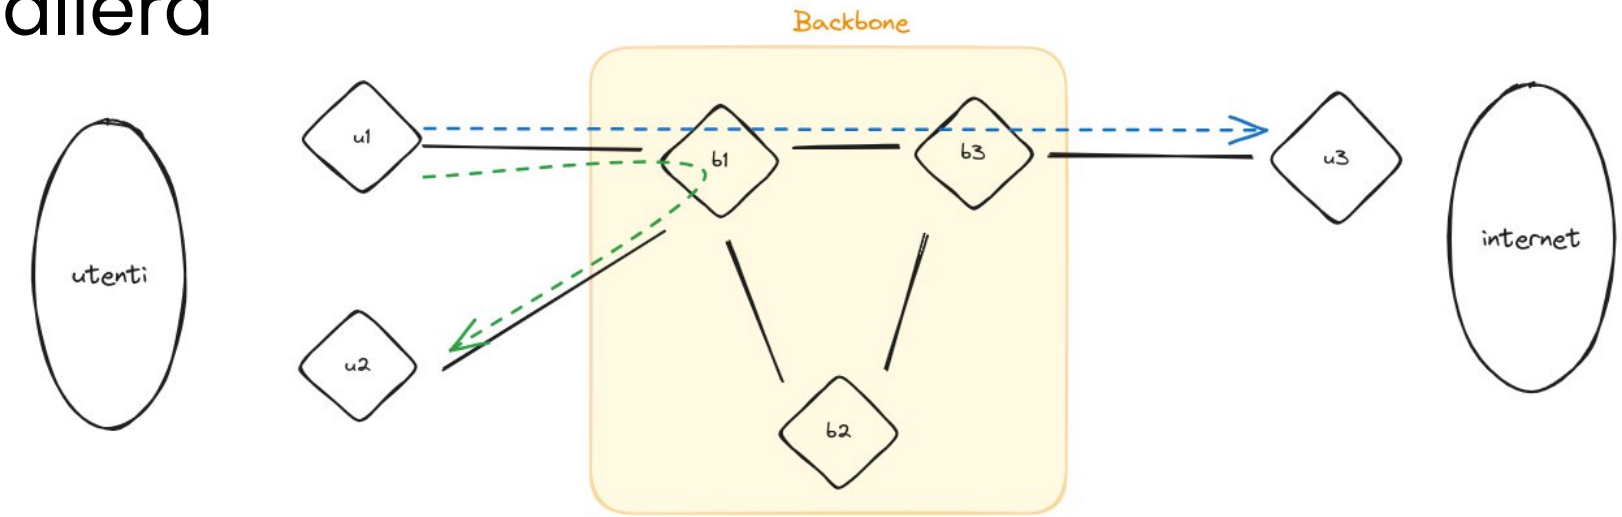

Casi d'uso

- **Traffico interno a GARR**
  - Esclusione traffico di reti non GARR

- **Ricerca del traffico delle sedi GARR e gli ASN di ricerca per la pianificazione dei link di GEANT**
  - Query per individuare gli ASN di ricerca (sui link GEANT)
  - Query entry relative ai soli ASN individuati

Dati aggregati per singolo ente o per categorie

Casi d'uso: per categorie di enti

dati 15gg

Traffico comunita' GARR ricevuto

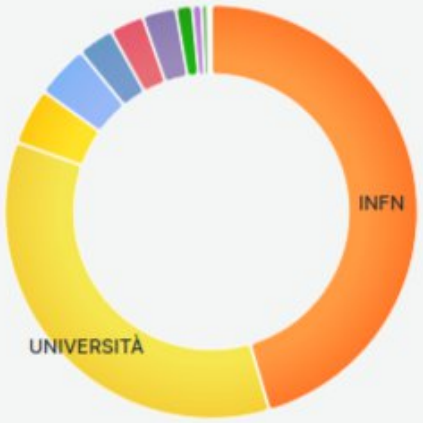

	Value
INFN	822 TB
UNIVERSITÀ	633 TB
CNR	79.0 TB
ENEA	74.3 TB
UNIVERSITÀ NON STATALI	49.8 TB
ENTI DI RICERCA SCIENTIFICA E TECNOLOGICA	47.2 TB
CONSORZI UNIVERSITARI	46.4 TB
ISTITUTI INTERNAZIONALI	21.7 TB
INAF	11.6 TB
ASI	7.45 TB
ISTITUTI RICERCA BIOMEDICA	1.29 TB
MUSEI SCIENTIFICI	778 GB

Traffico comunita' GARR inviato

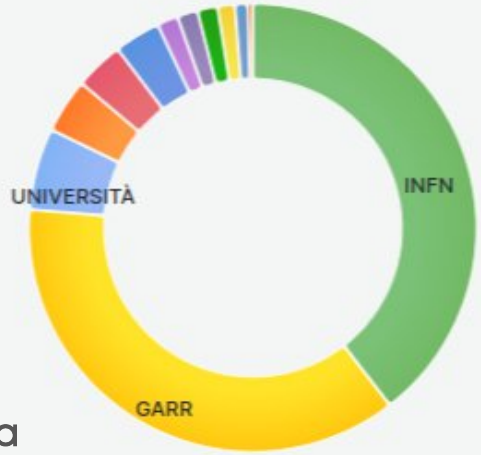

Grafana

	Value
INFN	760 TB
GARR	707 TB
UNIVERSITÀ	112 TB
ISTITUTI INTERNAZIONALI	74.5 TB
CONSORZI UNIVERSITARI	66.5 TB
ENEA	62.9 TB
INGV	27.6 TB
UNIVERSITÀ NON STATALI	27.0 TB
CNR	26.6 TB
ENTI DI RICERCA SCIENTIFICA E TECNOLOGICA	22.5 TB
ASI	16.6 TB
INAF	5.04 TB

Top traffico interno a GARR

Casi d'uso: per categorie di enti

dati 15gg

Traffico di ricerca ricevuto

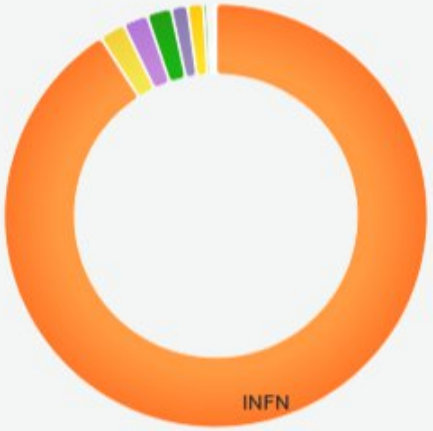

	Value
INFN	3.80 PB
UNIVERSITÀ	79.4 TB
INAF	75.8 TB
ISTITUTI INTERNAZIONALI	75.5 TB
CONSORZI UNIVERSITARI	48.8 TB
CNR	47.7 TB
ASI	9.40 TB
ENEA	4.05 TB
ENTI DI RICERCA SCIENTIFICA E TECNOLOGICA	3.88 TB
ISTITUTI RICERCA BIOMEDICA	1.65 TB
UNIVERSITÀ NON STATALI	375 GB
MUSEI SCIENTIFICI	2.91 GB

Traffico di ricerca inviato

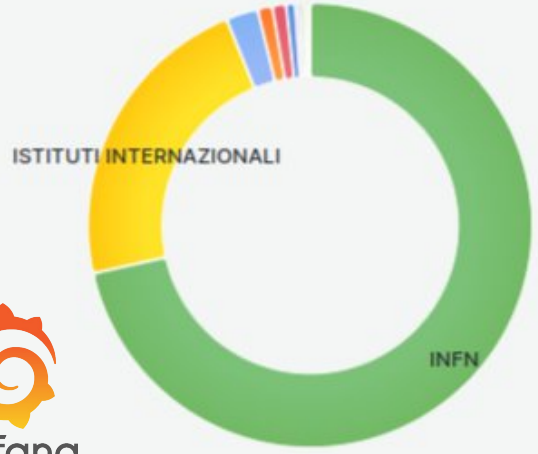

	Value
INFN	2.06 PB
ISTITUTI INTERNAZIONALI	642 TB
CONSORZI UNIVERSITARI	64.7 TB
UNIVERSITÀ	28.7 TB
INAF	26.5 TB
ASI	16.4 TB
GARR	4.43 TB
ARCHIVI	3.76 TB
CNR	3.68 TB
ENTI DI RICERCA SCIENTIFICA E TECNOLOGICA	3.48 TB
ISTITUTI RICERCA BIOMEDICA	1.66 TB
ENEA	761 GB

Grafana

Top traffico di ricerca con AS esterni

Casi d'uso: Singola Sede

dati 15gg

INGV - Grottaminarda (AV)

Top traffico interno comunita' GARR

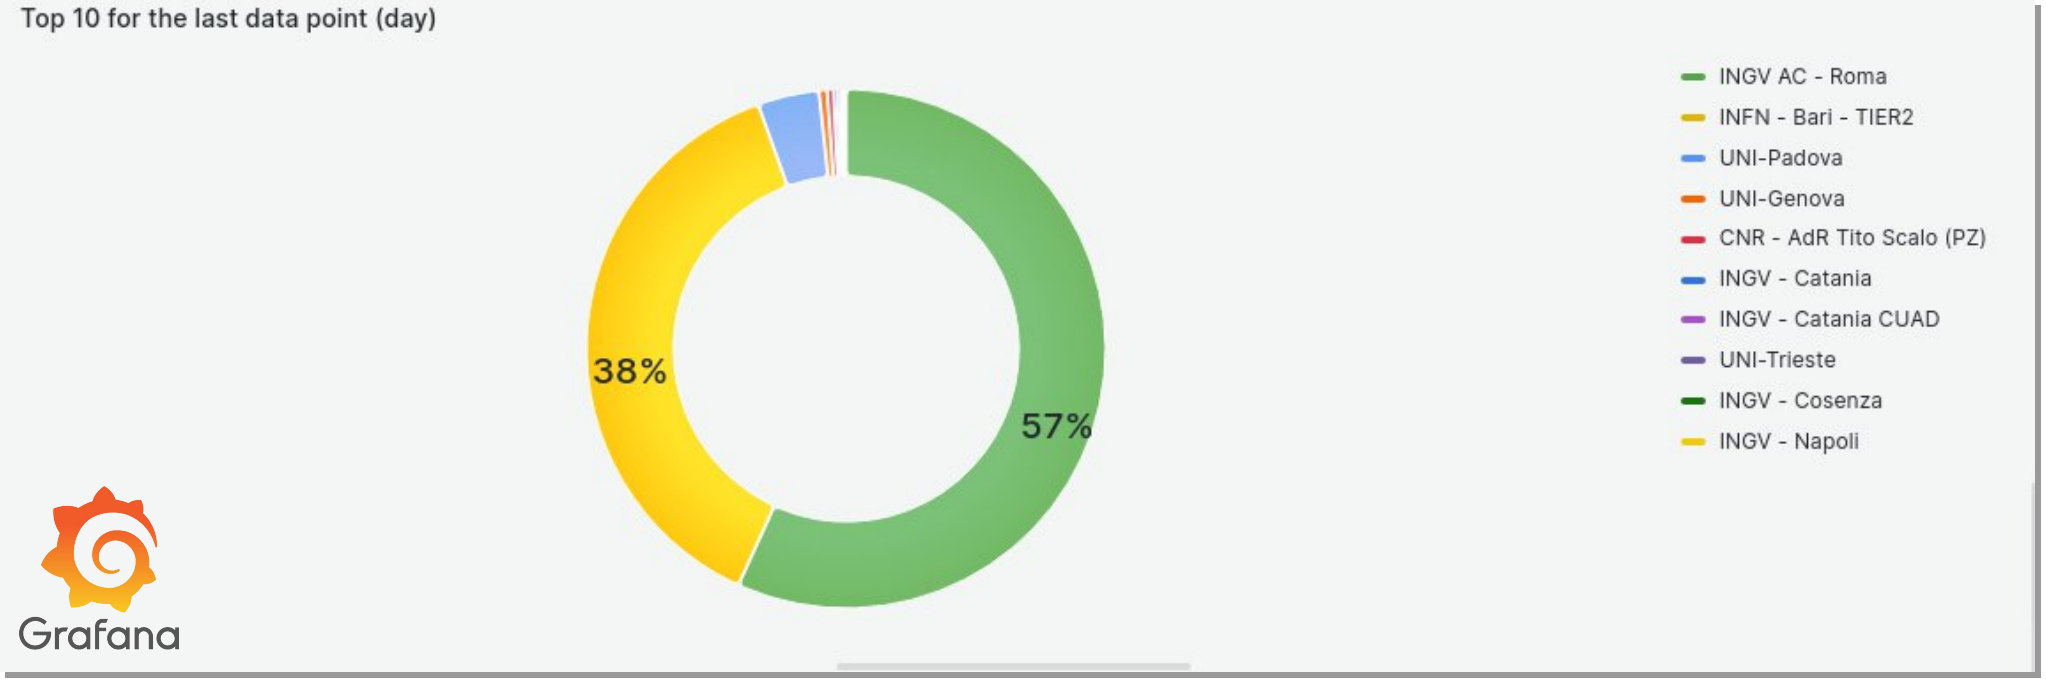

Traffico per tipologia

Name	research_bytes / Total Traffic	internal_bytes / Total Traffic	other_bytes / Total Traffic
INGV - Grottaminarda (AV)	28.5%	38.3%	33.2%

Conclusioni e sviluppi futuri

- Stabilizzare il sistema in ottica di produzione e data retention
- Nuove possibilita’ di sviluppo
- Porting dei servizi basati su nfdump

Fine

giovanni.cesaroni@garr.it
nino.ciurleo@garr.it