

Nov-28-2018 第 31 屆 TWNIC OPM 暨第二屆 TWNOC 會議：

Machine Learning and AI for Networking

Machine Learning, Popular Model, iCAM, Historical and Predictive Analytics,
Visualizations

錢小山

首席技術顧問

思科大中華區數據中心架構事業部

二〇一八年十一月

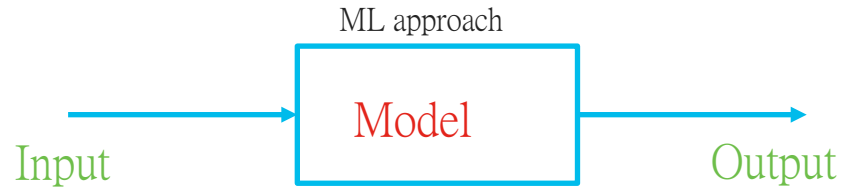



Machine Learning

Nov. 2018



Machine Learning Vs Classical approach



Supervised Learning

X is the known input
Y is the known output

$$Y = F(X)$$

Goal: Find the Function F

So that new outputs can be predicted using new inputs

Terminology :

X : training data

Y : output

F : Model

Finding F: Training the data

Features			Label
Bedrooms	Sq. feet	Neighborhood	Sale price
3	2000	Normaltown	\$250,000
2	800	Hipsterton	\$300,000
2	850	Normaltown	\$150,000
1	550	Normaltown	\$78,000
4	2000	Skid Row	\$150,000

training data

Test data



Unsupervised Learning

Only X inputs are known
Outputs Y, F are unknown

$$Y = F(X)$$

Goal:

- Group and interpret data based on input data X
- Find hidden structure
- Partition the data set



Popular Machine Learning Libraries





Popular Model

Nov. 2018



Linear Regression

Constant

Coefficient


$$y = b_0 + b_1 x_1$$

Dependent Variable

Independent Variable

Linear Regression Example

Experience Vs Salary



Polynomial Regression

Linear Regression

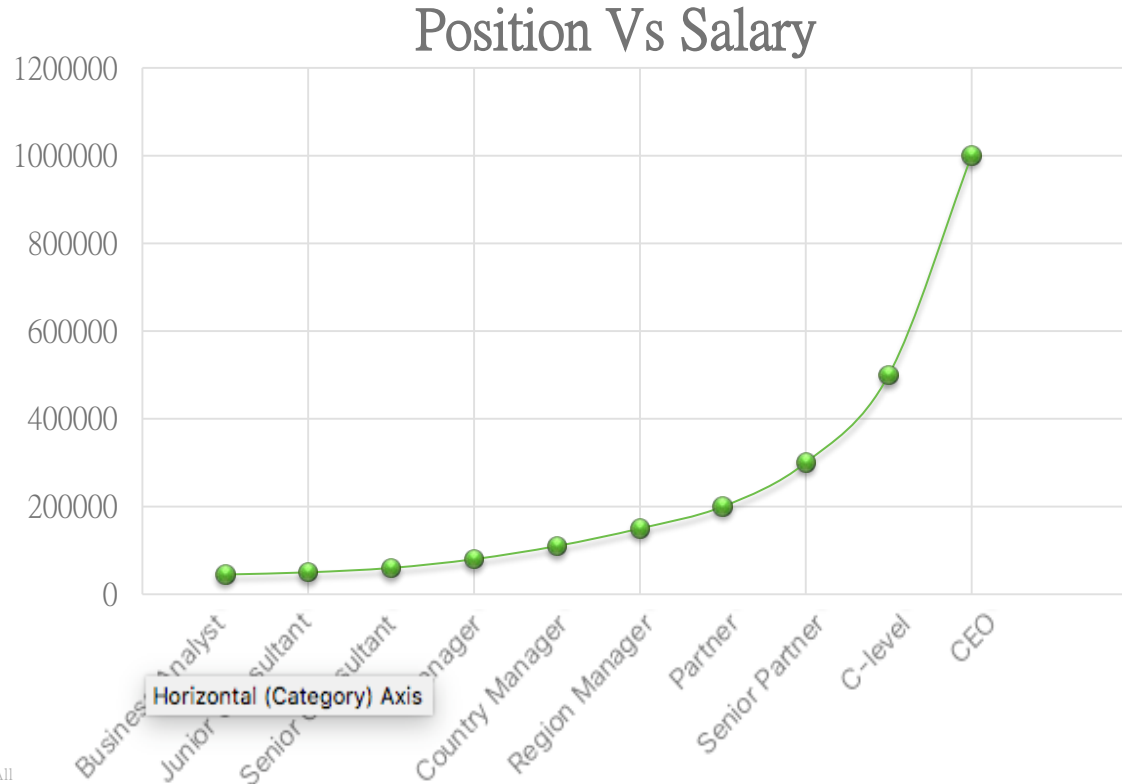
$$y = b_0 + b_1x_1$$

Multi-Linear Regression


$$y = b_0 + b_1x_1 + b_2x_2 + \dots + b_nx_n$$

Polynomial Regression $y = b_0 + b_1x_1 + b_2x_1^2 + \dots + b_nx_1^n$


Polynomial Regression Example




Decision Tree Intuition



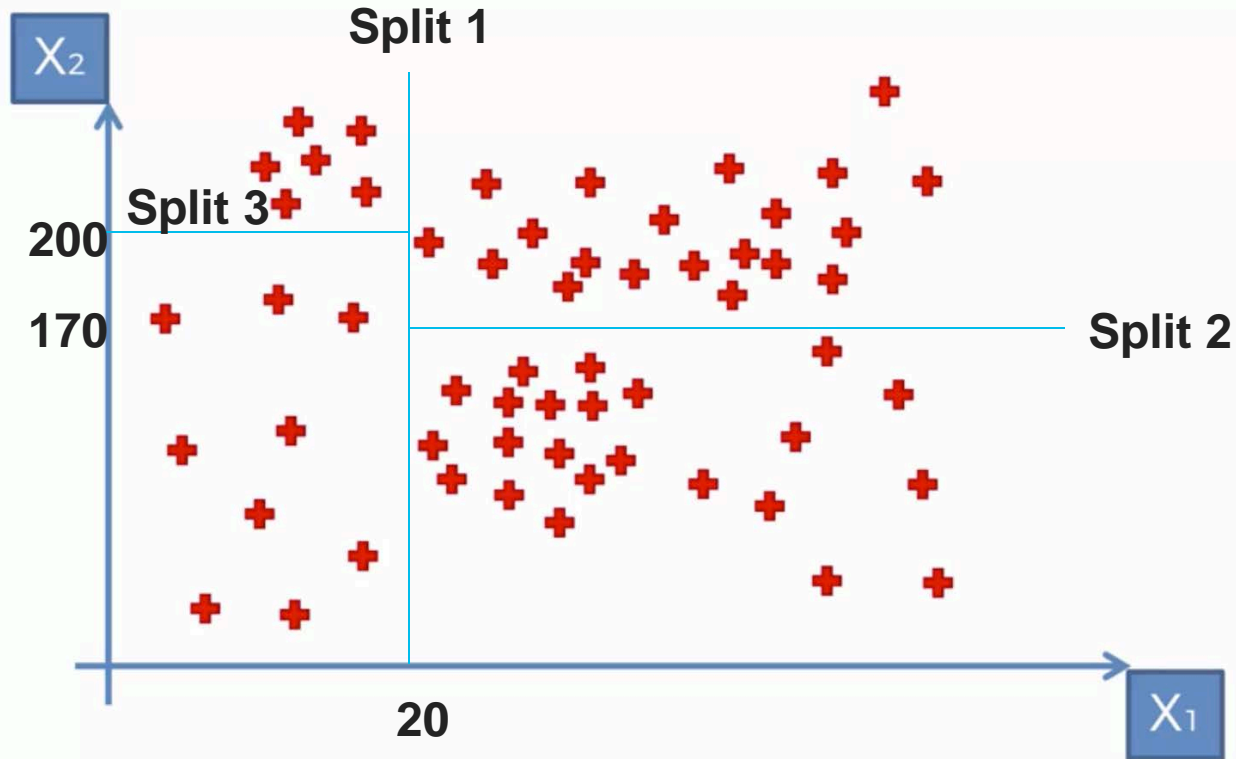
Decision Tree Intuition




Decision Tree Intuition



Decision Tree Intuition



Decision Tree Intuition



Sample Code: Linear Regression

```
#Importing the libraries
import numpy as np
import matplotlib.pyplot as plt
import pandas as pd

# Importing the dataset
dataset = pd.read_csv('Position_Salaries.csv')
X = dataset.iloc[:, 1:2].values
y = dataset.iloc[:, 2].values

# Fitting Linear Regression to the dataset
from sklearn.linear_model import LinearRegression
regressor = LinearRegression()
regressor.fit(X, y)

# Visualising the Linear Regression results
plt.scatter(X, y, color = 'red')
plt.plot(X, regressor.predict(X), color = 'blue')
plt.title('Truth or Bluff (Linear Regression)')
plt.xlabel('Position level')
plt.ylabel('Salary')
plt.show()

# Predicting a new result with Linear Regression
regressor.predict(6.5)
```

Sample Code: Decision Tree Regression

```
# Decision Tree Regression

# Importing the libraries
import numpy as np
import matplotlib.pyplot as plt
import pandas as pd

# Importing the dataset
dataset = pd.read_csv('Position_Salaries.csv')
X = dataset.iloc[:, 1:2].values
y = dataset.iloc[:, 2].values

# Fitting Decision Tree Regression to the dataset
from sklearn.tree import DecisionTreeRegressor
regressor = DecisionTreeRegressor(random_state = 0)
regressor.fit(X, y)

# Predicting a new result
y_pred = regressor.predict(6.5)

# Visualising the Decision Tree Regression results (higher resolution)
X_grid = np.arange(min(X), max(X), 0.01)
X_grid = X_grid.reshape((len(X_grid), 1))
plt.scatter(X, y, color = 'red')
plt.plot(X_grid, regressor.predict(X_grid), color = 'blue')
plt.title('Truth or Bluff (Decision Tree Regression)')
plt.xlabel('Position level')
plt.ylabel('Salary')
plt.show()
```



iCAM

Nov. 2018




Problem Statement

- Currently customers do not know about
- Historical trends of traffic: per-subnet, per application
- Future predictions about traffic: per-subnet, per-application
- Top heavy hitters of traffic
- Historical trends of hardware/software tables on the switches
- Future predictions about hardware/software table utilizations
- When is a good time to schedule downtime?
- Which hardware entries can be retired to create space for new ones?
- Are there any anomalies (quick visualization)?

Solution: iCAM

- iCAM = Intelligent Comprehensive Analytics and Machine Learning
- Machine Learning natively on the switches.
- Already Shipped on Nexus 9000 and Nexus 7000 Series of switches
- Ship Date: Dec 2016
- Software License
- Several Patents pending

iCAM



Contact : nxos-icam@cisco.com




iCAM Overview:

- Machine Learning, natively on the switch/router
- Security access control analytics
- Internal hardware tables usage analytics
- Top/bottom heavy hitters
- Anomaly visualization
- Build apps on top of iCAM
- Historical Analytics
- Predictive Analysis
- Streaming telemetry

Benefits:



- Order of magnitude OPEX savings : Predictive analytics and historical analytics
- Order of magnitude CAPEX savings : Natively on the switch/router: Wiring, Power, Rackspace and Cost savings
- Scalability : Multi-Terabits/s
- Security: Anomaly visualization

Interactive Graph/GUI



Thick end is the source IP/subnet.
Thin end is the dest IP/subnet.
Thickness represents the % traffic.

Interactive Graph/GUI





Historical and Predictive Analytics

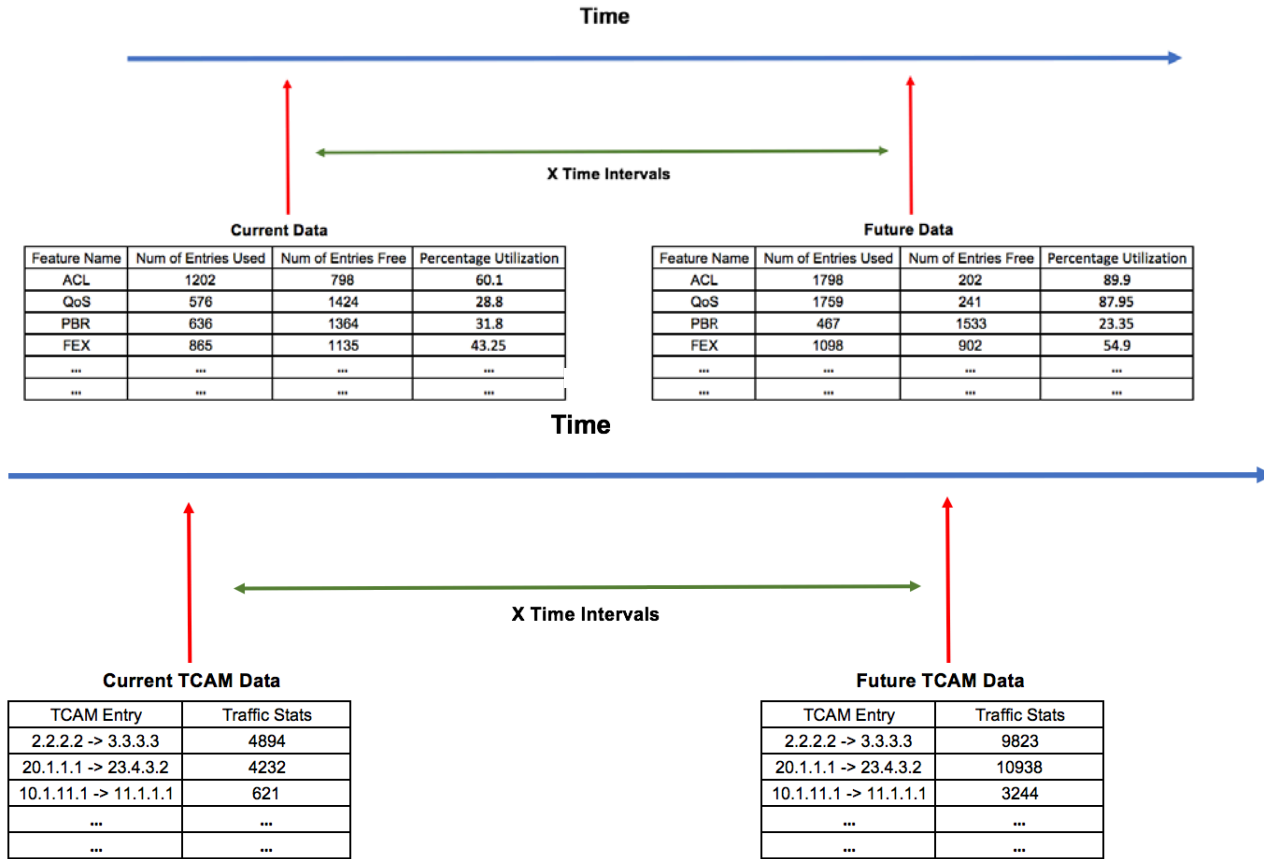
Nov. 2018



Show ML Predictions

- `show icam prediction entries {acl | multicast} module module-number inst instance-number year month day HH:MM:SS [top x]`
- `show icam prediction resource {acl_tcam | fib_cam | l2_table} module module-number inst instance-number year month day HH:MM:SS`

iCAM Predictions



Example – Multicast Traffic Prediction

```
switch# show icam prediction entries multicast module 5 2018 July 10 10:00:00
```

```
=====
```

```
Multicast Entries Prediction (Mod 5)
```

```
-----
```

VDC_ID	TABLE_ID	Source/Mask	Group/Mask	RPF	Stats	Prediction
3	32768	0.0.0.0/0	224.0.0.0/24		3198	3364
2	32768	0.0.0.0/0	224.0.0.0/24		7636	8421
2	3	200.1.1.2/32	225.1.1.229/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.125/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.126/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.127/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.128/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.129/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.130/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.131/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.132/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.133/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.134/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.228/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.202/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.35/32	Ethernet5/1	2554	3192
2	3	200.1.1.2/32	225.1.1.36/32	Ethernet5/1	2554	3192
2	3	200.1.1.2/32	225.1.1.232/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.231/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.201/32	Ethernet5/1	1	1
2	3	200.1.1.2/32	225.1.1.41/32	Ethernet5/1	2554	3192
2	3	200.1.1.2/32	225.1.1.42/32	Ethernet5/1	2554	3192
2	3	200.1.1.2/32	225.1.1.43/32	Ethernet5/1	2554	3192
2	3	200.1.1.2/32	225.1.1.44/32	Ethernet5/1	2554	3192

Example - ACL Resource Prediction

```
switch# show icam prediction resource acl_tcam module 5 inst 1 2018 July 10 10:00:00
```

```
-----  
Feature Hardware Resource Prediction (Mod 5,Inst 1)  
-----
```

Feature	Direction	TCAM#	BANK#	Feature_Entries	Free_Entries	Percent_Util
FEX Control	ingress	1	0	5	32743	0.00
CoPP	ingress	1	1	442	32306	1.00

```
=====
```

```
ACL TCAM Resource Prediction (Mod 5,Inst 1)
```

```
-----
```

	Used	Free	Percent_Util
Tcam 1 Bank 1	462	32306	1.40
Tcam 1 Bank 0	25	32743	0.07
Tcam 0 Bank 1	20	32748	0.06
Tcam 0 Bank 0	20	32748	0.06

Example – FIB Resource Prediction

```
switch# show icam prediction resource fib_tcam module 5 inst 0 2018 July 10 10:00:00
```

```
=====
FIB TCAM Resource Prediction (Mod 5, Inst 0)
-----
```

Type	logical	physical	Percent_Util
FCMPLS	0	0	0.00
IPV4 unicast	132	132	0.00
DIAG_80	0	0	0.00
EOM Peer	0	0	0.00
MPLS	0	0	0.00
IPV6 multicast	5	20	0.00
IPV6 LinkLocal	1	2	0.00
FCOE	0	0	0.00
MPLS VPN	0	0	0.00
IPV4 multicast	11273	11273	1.00
IPV6 unicast	2	4	0.00




Visualizations of iCAM

Nov. 2018




Interactive Graph/GUI

ic



Anomaly Visualization



One specific TCAM entry has too much traffic

iCAM Supported Platforms/Software Release

Platform



Nexus 7000/7700 Series



Nexus 9000 Series

Version

NX-OS 8.0(1)

NX-OS 7.0(3)I7(1)

License

Enhanced Layer 2

Network Services

CLI - Configuration

- feature icam
- icam monitor interval *interval-hours* num_intervals *number-of-intervals*
- icam monitor resource {acl_tcam | fib_tcam | l2_table} module *module-number* inst *instance-number*
- icam monitor entries {acl | multicast} module *module-number* inst *instance-number*

Notes

- The monitor commands are optional

CLI – Show Analytics - 2

Sorting and Filtering the output.

- **sort** Sorts the entries based on filter and sort-order.
- **filter** Filters the entries based on the feature name. Any feature name as a string such as " RACL" , "QoS" , etc., can be used to filter the output.
- **exact** Exact keyword. Eg, "QoS CoPP"
- **sort-order** Sorts the entries in either ascending or descending order.
- **top** Displays the top x% iCAM entries for a feature. Values range from 1 to 100.
- **history** Provide historical analytics.

Example:

```
show icam entries acl module 1 instance 0 sort filter RACL sort-order ascending top 10
```

iCAM Resources History - Example

Stats snapshots in last X intervals

```
n35(config)# sh icam resource fib_tcam module 5 inst 0 history 3
```

```
=====
FIB TCAM Resource Utilization (Mod 5, Inst 0)
-----
```

Type	logical	physical	Percent_Util	Timestamp (UTC)
IPV4 unicast	16	16	0.00	2017-09-12 06:06:53
	16	16	0.00	2017-09-12 06:16:55
	16	16	0.00	2017-09-12 06:26:37
	16	16	0.00	2017-09-12 06:46:23
	16	16	0.00	2017-09-12 06:56:59
DIAG_80	0	0	0.00	2017-09-12 06:06:53
	0	0	0.00	2017-09-12 06:16:55
	0	0	0.00	2017-09-12 06:26:37
	0	0	0.00	2017-09-12 06:46:23
	0	0	0.00	2017-09-12 06:56:59
IPV4 multicast	6	6	0.00	2017-09-12 06:06:53
	45	45	0.10	2017-09-12 06:16:55
	45	45	0.10	2017-09-12 06:26:37
	62	62	0.17	2017-09-12 06:46:23
	62	62	0.17	2017-09-12 06:56:59
MPLS	0	0	0.00	2017-09-12 06:06:53
	0	0	0.00	2017-09-12 06:16:55
	0	0	0.00	2017-09-12 06:26:37
	0	0	0.00	2017-09-12 06:46:23
	0	0	0.00	2017-09-12 06:56:59

iCAM Entries History - Example

```
switch# show icam entries acl module 5 inst 0 history 5
```

Cumulative Stats in the last X intervals




```
=====
```

```
TCAM Entries (Mod 5,Inst 0): Cumulative stats for last 5 intervals
```

```
-----
```

Feature	Pkt Type	Source IP/Mask	Dest IP/Mask	Action	ifindex	Stats	Rate(pps)
FEX	IPv4	ip 0.0.0.0/0	0.0.0.0/0	Redirect	0x15090000	0	0
FEX	IPv6	ip 0x00000000000000000000000000000000/0	0x00000000000000000000000000000000	Redirect	0x15090000	0	0
FEX	MAC	0000.0000.0000	0000.0000.0000	Redirect	0x15090000	0	0
FEX	ARP	arp-rarp/all ip 0.0.0.0/0	0.0.0.0/0	Redirect	0x15090000	0	0
RACL	IPv4	ip 2.2.2.1/32	1.1.1.1/32	Permit	0x1a200000	23423	945
RACL	IPv4	ip 2.2.2.2/32	1.1.1.2/32	Permit	0x1a200000	23946237	718353
RACL	IPv4	ip 2.2.2.3/32	1.1.1.3/32	Permit	0x1a200000	83675	585
RACL	IPv4	ip 2.2.2.4/32	1.1.1.4/32	Permit	0x1a200000	0	0
RACL	IPv4	ip 2.2.2.5/32	1.1.1.5/32	Permit	0x1a200000	9693487	45986
RACL	IPv4	ip 2.2.2.9/32	1.1.1.9/32	Permit	0x1a200000	9693487	45986
RACL	IPv4	ip 0.0.0.0/0	0.0.0.0/0	Deny	0x1a200000	0	0
QoS	COPP	tcp 0.0.0.0/0	0.0.0.0/0	QoS	0x0	0	0
QoS	COPP	udp 0.0.0.0/0	224.0.0.0/24	QoS	0x0	0	0

Average Rate in the last X intervals

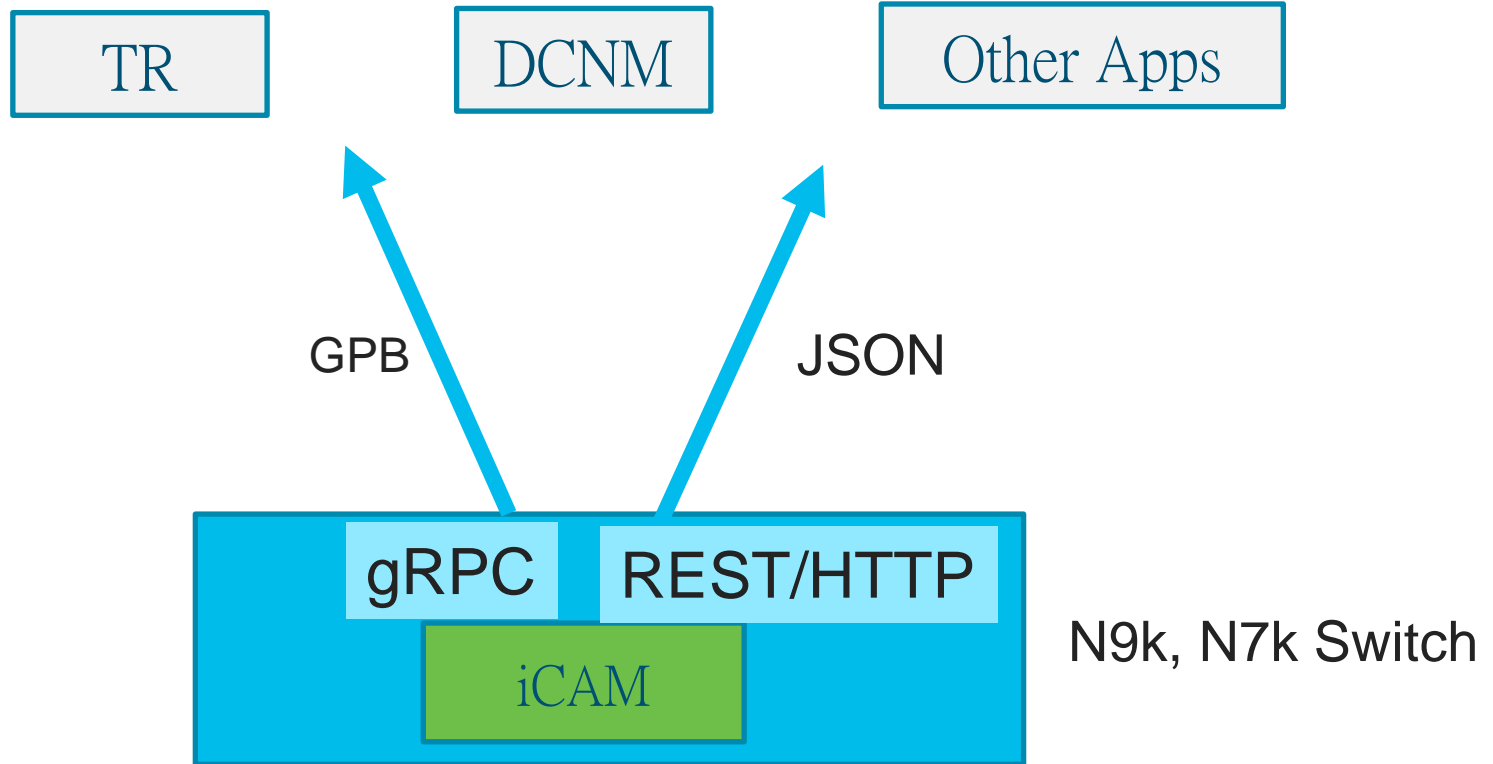


Example - Layer 2 MAC Table

```
switch# sh icam resource l2_table module 1 inst 0
```

```
=====
L2 Table Resource Utilization (Mod 1,Inst 0)
-----
  Total_entries   Used_entries   Percent_Util           Timestamp (UTC)
-----
           98304             341             0.34           2017-11-07 01:21:20
```

Streaming Telemetry with iCAM



- Dashboard
- Topology
- Inventory

- Monitor

- Configure

- Administration

System tools

	Name	Description
1	iCAM Traffic	Provides information about traffic exchange per switch based on CAM table configuration
2	iCAM Usage	Provides information about CAM table configuration per switch

Awards

- Best of Interop ITX 2017: [Finalist in Security Category](#)
- Best of Interop ITX 2017: [Finalist in DevOps Category](#)
- Light Reading Awards: Finalist in [Outstanding Communications Technology Vision](#)
- Best in Biz Awards: [Most innovative Product of the Year](#)

More media links

- Chuck Robbins, Cisco CEO tweet:
 - <https://twitter.com/ChuckRobbins/status/861961430559371264>
 - <https://twitter.com/ChuckRobbins/status/869319371398586368>
- Frank Palumbo, Cisco SVP, Worldwide Sales tweet:
 - <https://twitter.com/fpalumbo/status/861738133196480512>
 - <https://twitter.com/fpalumbo/status/870346475439566848>
- Network Computing about our innovations:
 - <http://www.networkcomputing.com/cloud-infrastructure/meet-best-interop-itx-2017-finalists/1224184367>

References

- [iCAM config guide](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_011010.html) https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_7x_chapter_011010.html
- <http://blogs.cisco.com/datacenter/icam>
- nxos-icam@cisco.com

Thank You

