Lancaster University

SDN Heading North: Towards a Declarative Intent-based Northbound Interface

Shiyam Alalmaei, Matthew Broadbent, Yehia Elkhatib, Mehdi Bezahaf, and Nicholas Race

School of Computing and Communications, Lancaster University, United Kingdom

Email: {s.alalmaei, m.broadbent, y.elkhatib, mehdi.bezahaf, n.race}@lancaster.ac.uk

Motivation

Lancaster University

- **Why is the NBI important?**

- The success of SDN relies on application developers leveraging its capabilities.

- Hence, the Northbound Interface (NBI) is the key enabler for the realization of the ultimate SDN promise.

- Although Intent-based NBIs are gaining a lot of attention, their development remains in its infancy.

Outline

Lancaster University

- Intent-Based NBI Challenges
- Related Work Limitations
- Proposed Intent-Based NBI framework and expressions
- A proof-of-concept cloud CDN use case of a caching intent
- Intent Refinement
- Conclusion & Future Work

Intents vs. Policies

	Similarities	Differences
Policies[1]	Network and system abstraction	• *Prescriptive* rules • Specify set of Event-Condition-Action (ECA) rules and determine precisely what to do under different circumstances and triggers (how to do?) • Used by system experts who can articulate the set of rules • System behavior defined proactively
Intents		• *Declarative* expressions • Express desired outcome (what to do?) • Can be used by different users including service consumers and non-experts without enumerating rules • Could be a learning reactive system

Intent-Based NBIs Challenges

Lancaster University

- Service consumers require ***declarative*** rather than ***prescriptive*** intent expressions.

- Translating service-oriented intents to system operations needs ***intermediate interpretation*** as policies.

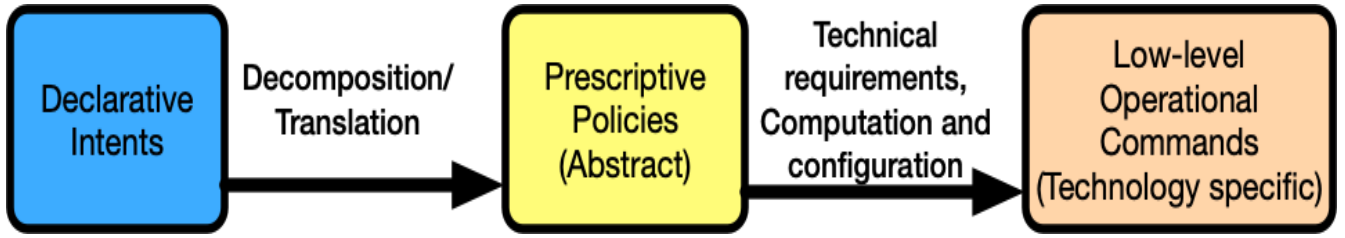

- Intent-based NBI must be ***platform-independent*** and ***extensible***

Related Work Limitations

Lancaster University

- In general, current Intent-based NBIs are limited, ad-hoc and vendor-specific.
- They do not allow expressing declarative intents that handle other requirements beyond the network-level.
- Moreover, most current works don’t provide the tools to create new intents and map them to lower-level policies.

Intent-Based NBI Framework

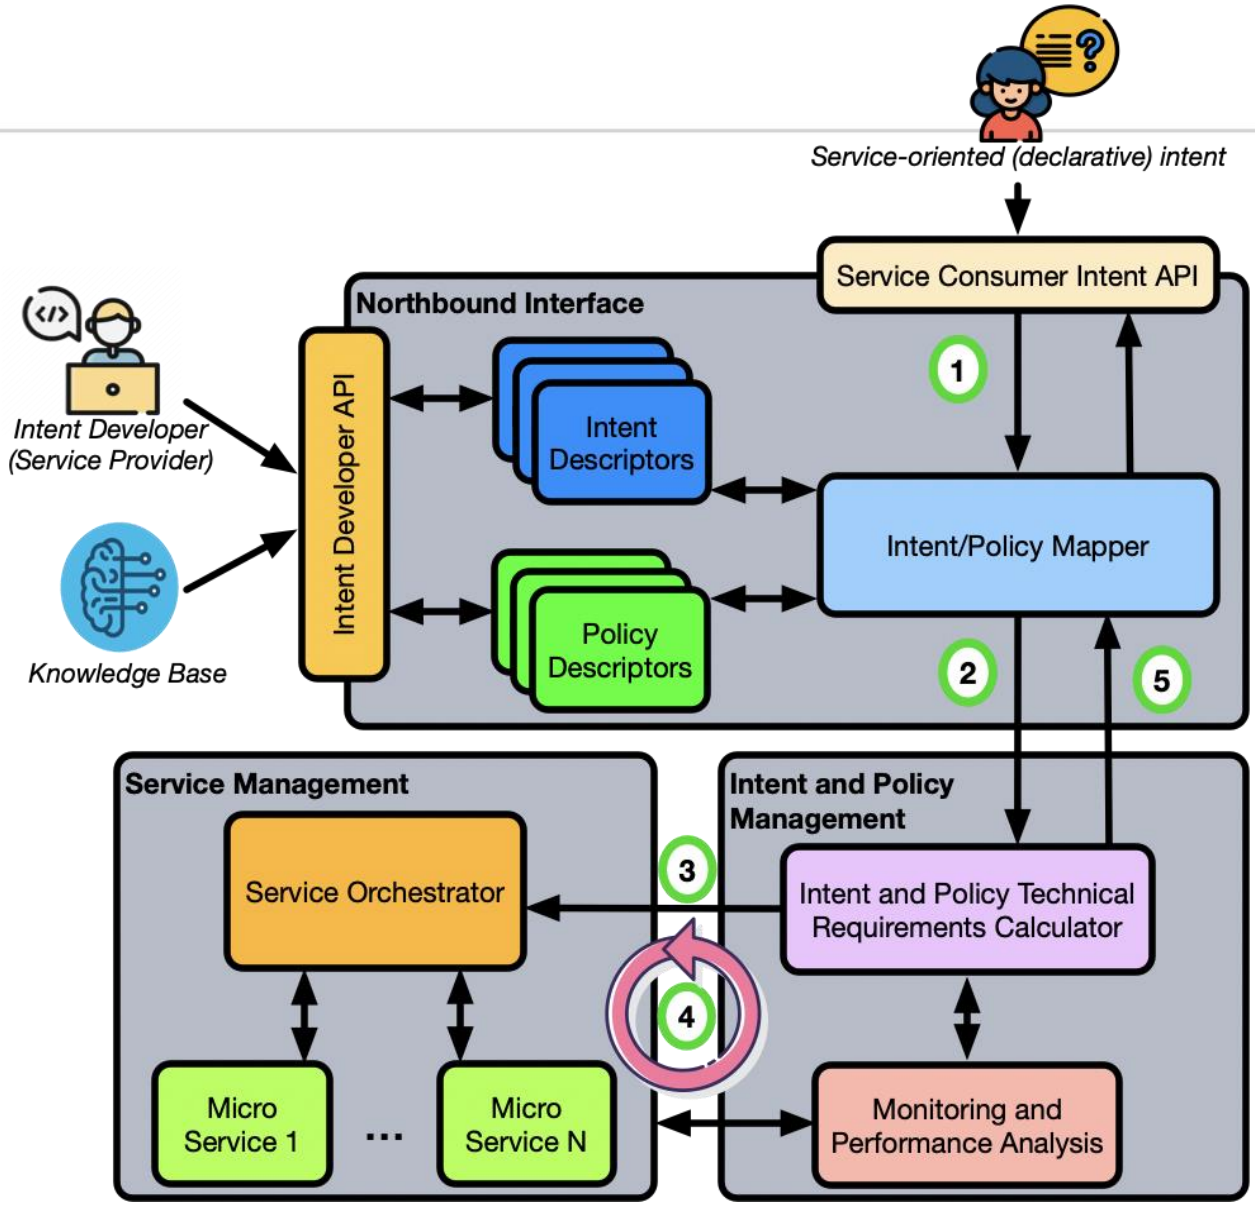

Intent-Based NBI Syntax

Declarative Service-Oriented Intent Syntax

<SERVICE><RESOURCES><CONJUNCTION><TARGET>

Prescriptive Policies Syntax

<SERVICE><RESOURCES>
<POLICY 1><OPERATOR><POLICY 2><OPERATOR>...
{<POLICY k_1><OPERATOR><POLICY k_2><OPERATOR>...<POLICY k_m>}
<OPERATOR>...<POLICY n>

<CONDITIONS><ACTIONS><CONSTRAINTS>[<PRIORITY>]

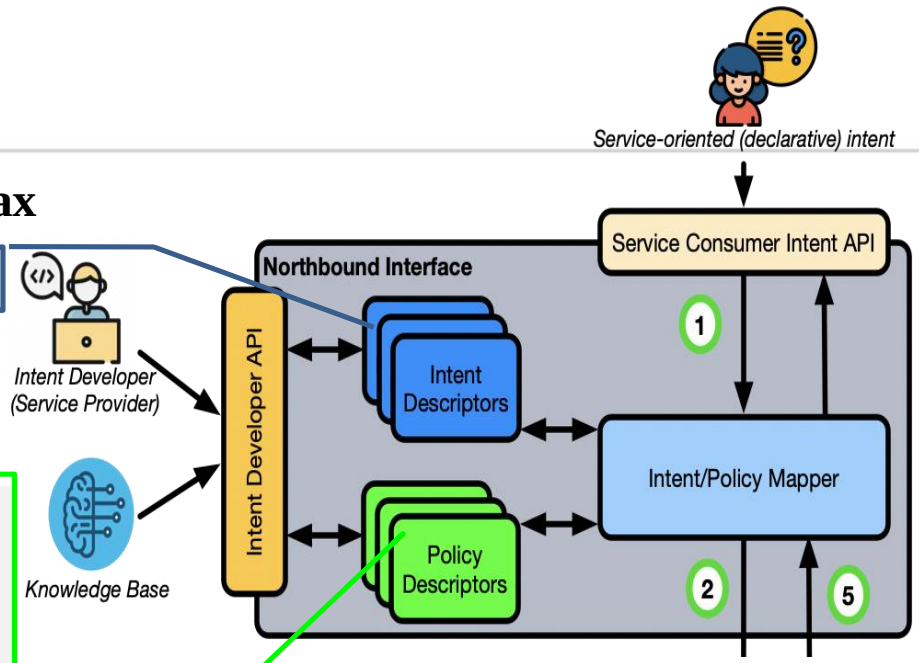

Intent/Policy Mapping

Tag	Basic Expression
<SERVICE>	Caching
<RESOURCES>	contents to be cached
<CONJUNCTION>	that can handle, that can meet, etc.
<TARGET>	<WORKLOAD>
<WORKLOAD>	<NUMBER> <UNIT> or <ADJECTIVE> <UNIT> or <ADJECTIVE> "workload"
<NUMBER>	numeric values that can represent the workload
<UNIT>	GB/min, requests/sec, etc.
<ADJECTIVE>	max, min, dynamic, high, medium, etc.
<CONDITIONS>	new caching request, max threshold exceeded, ...
<ACTIONS>	allocate cache servers(), scale out(), ...
<CONSTRAINTS>	with max storage, with least latency, ...
<PRIORITY>	optional indicator of policy priority
<OPERATOR>	Policy will be executed in parallel / sequential way

Cloud CDN (CCDN) Use Case

In our solution (opposed to Cloud CDNs today), Content Providers can express their high-level declarative intents targets (e.g., expected caching workload) which leads to better caching and resource management decisions with respect to the intent's target (well-, under-, and over-) estimation.

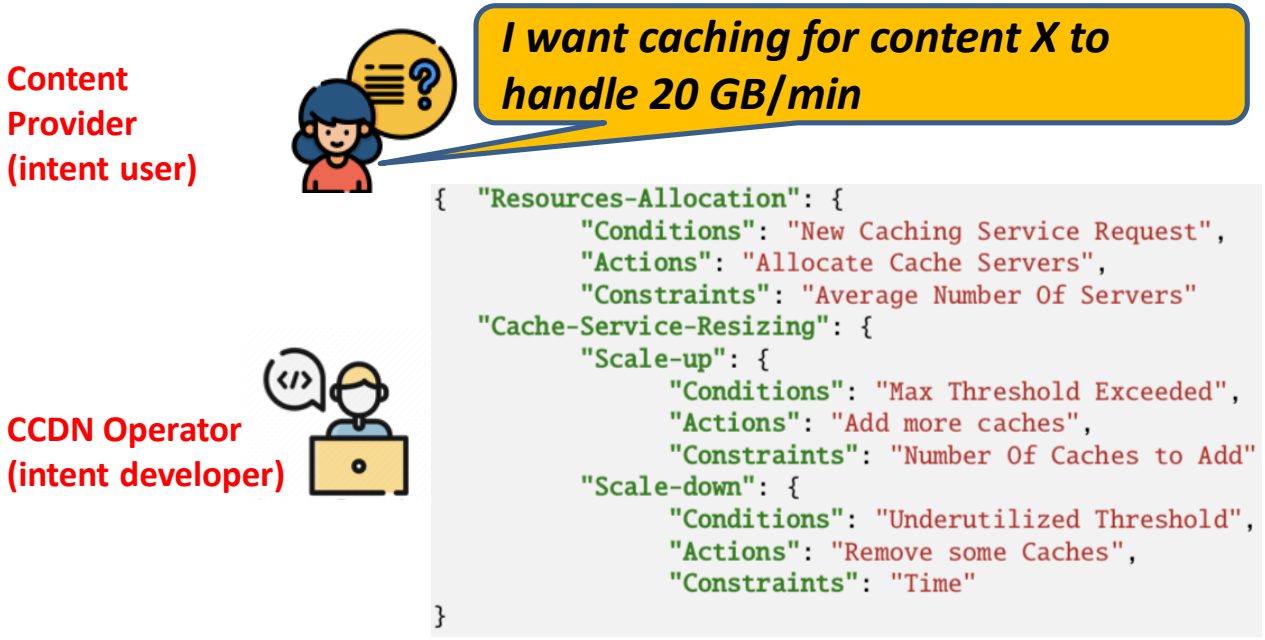

Cloud CDN (CCDN) Experimental Results

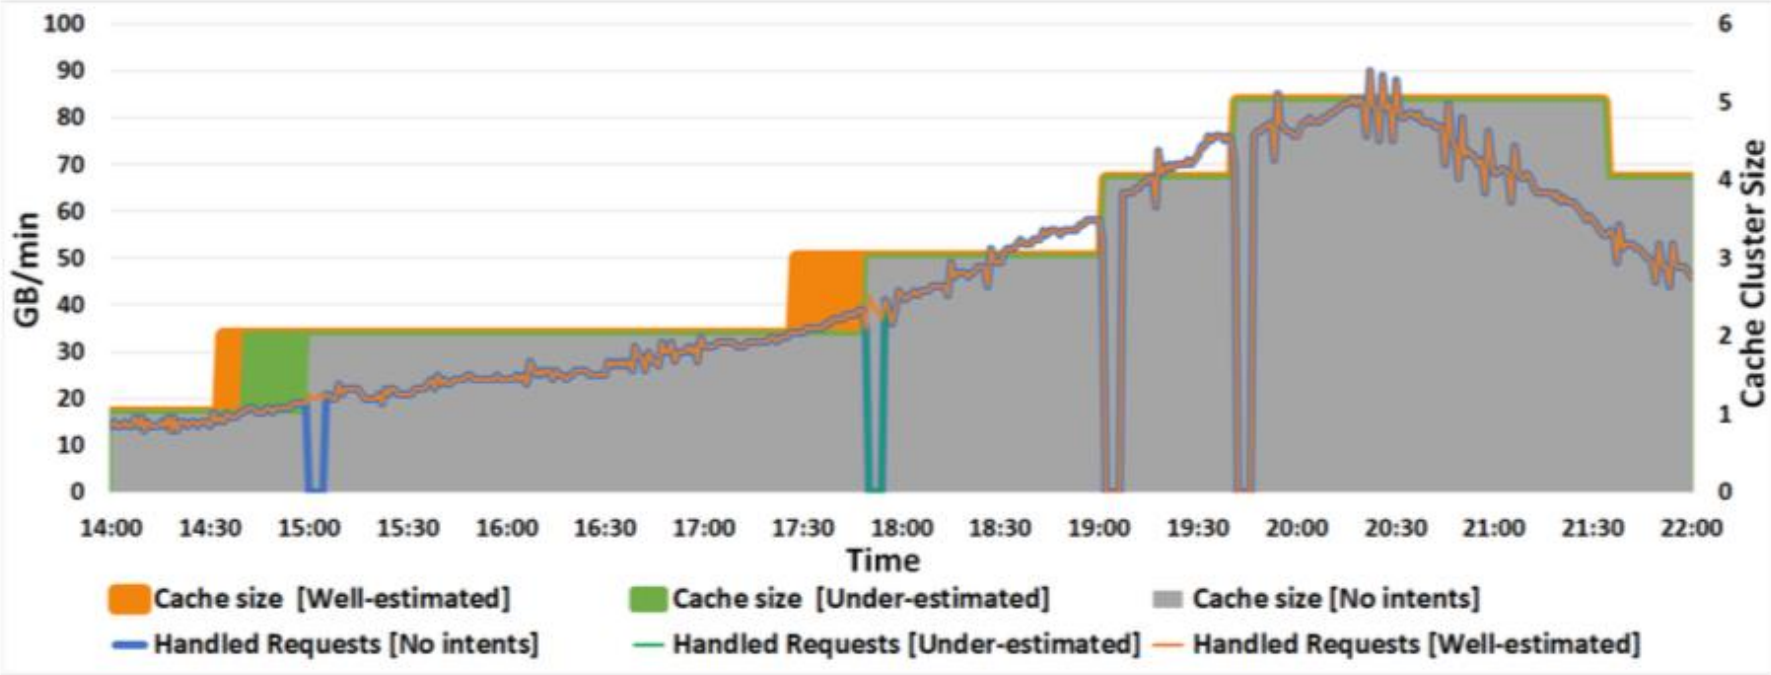

Cache cluster size and handled data rate in a reactive traditional CCDN (without intents) and an Intent-based one.

Intent Refinement

Lancaster University

To maximize the handled workload, the intent system would refine intent policies based on the previous demands and cluster size analysis. A way to achieve this is to scale out/in **proactively**. Two approaches can be taken:

- ***conservative*** method - aims to minimize the cost of deploying extra caches by sacrificing some unhandled requests;

- ***greedy*** method - handles more requests at the expense of additional deployment cost.

Intent Refinement Results

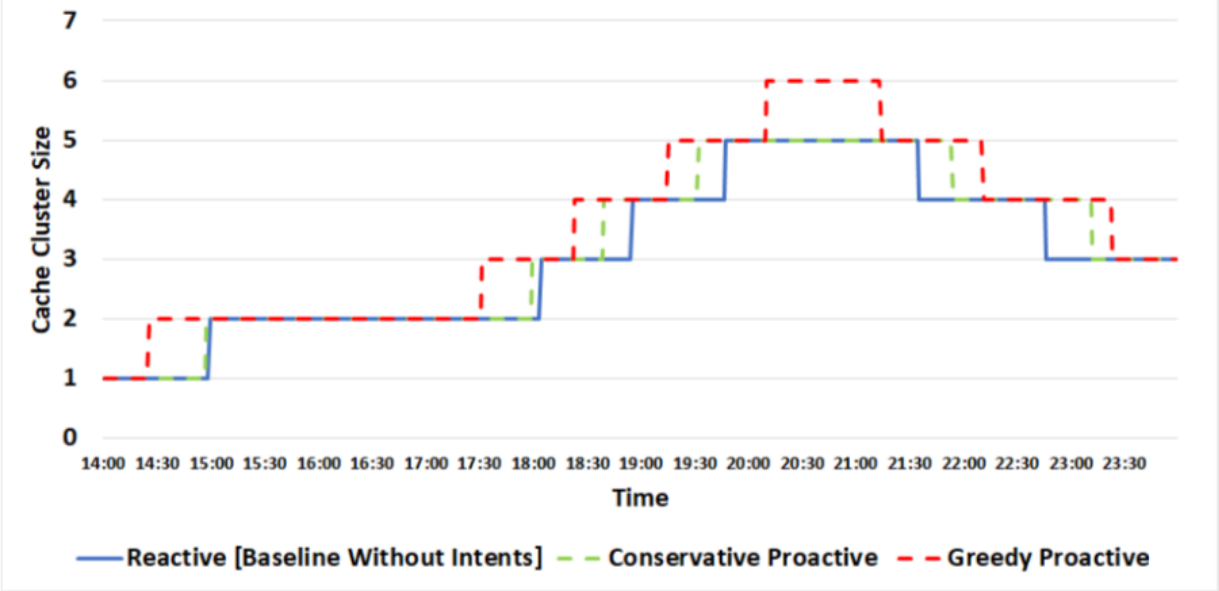

Cache cluster resizing (Conservative vs. Greedy policy)

	Traditional (No intent)	Intent-Based, conservative	Intent-Based, Greedy	Oracle
Unhandled Requests	3.17% ≃ 967 GB	0.96% ≃ 295 GB	0% ≃ 0 GB	0% ≃ 0 GB
Cost (minutes)	0	82	286	20

Intent Refinement Results

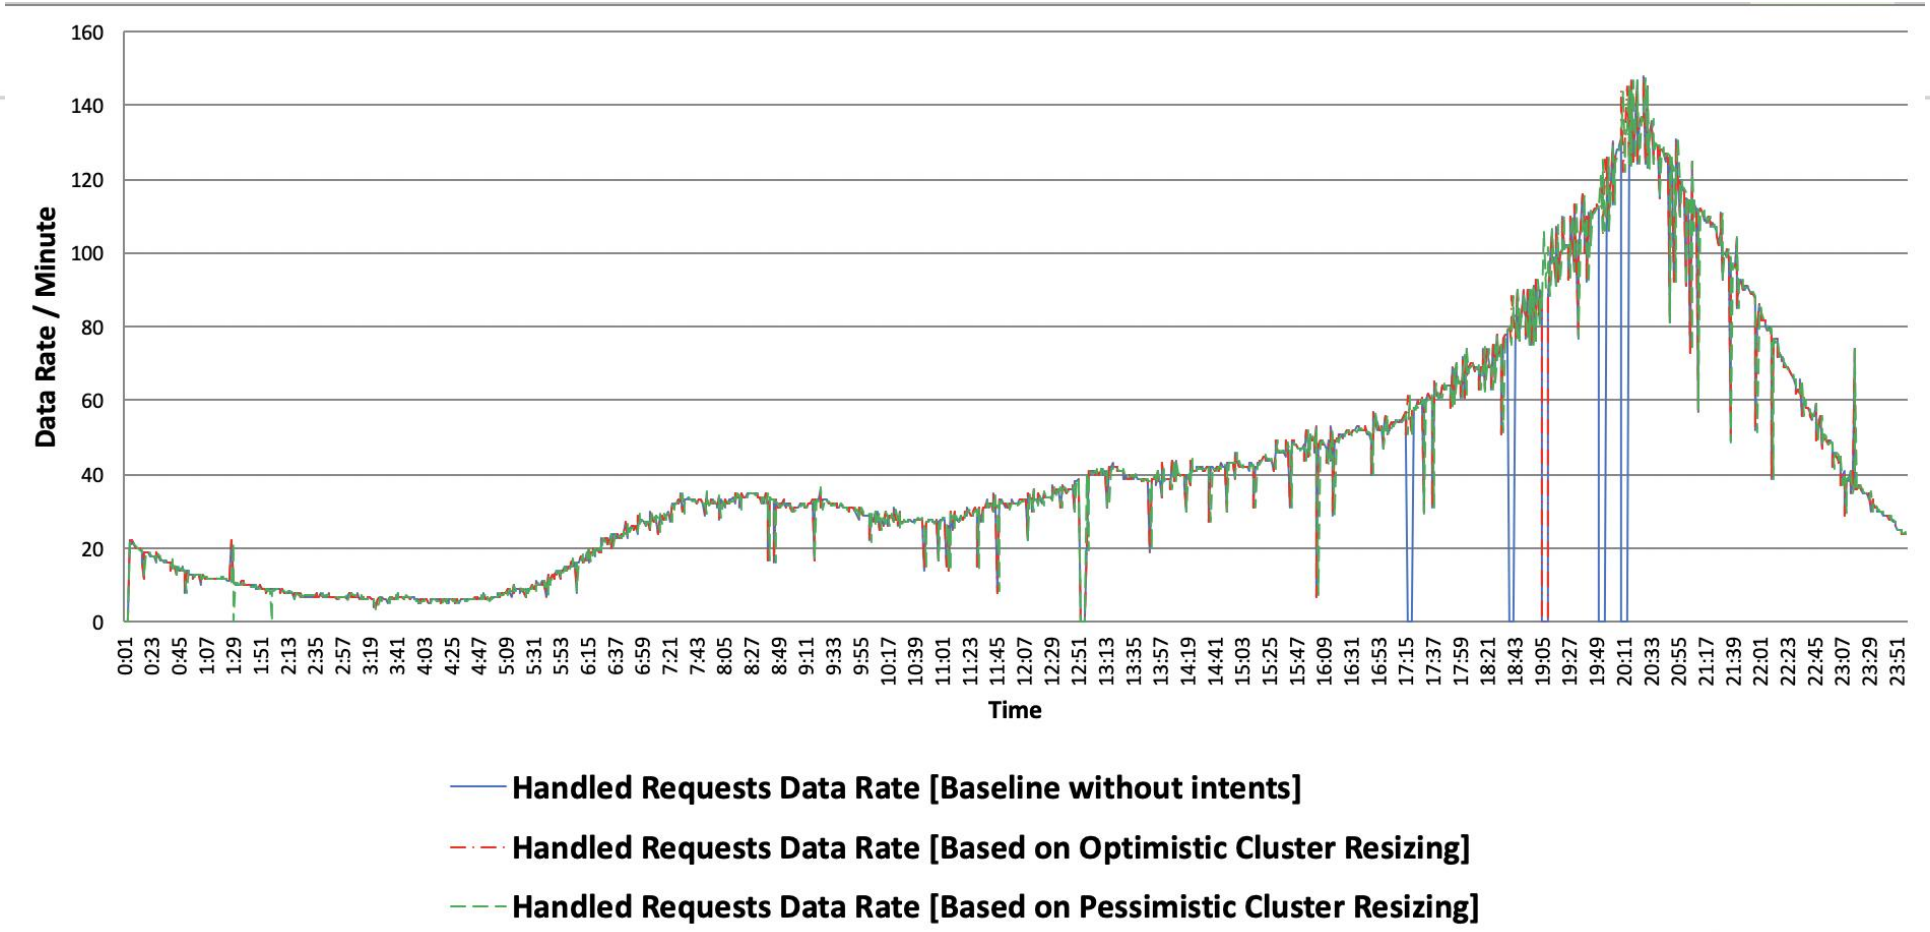

Handled Requests (Conservative vs. Greedy policy)

Conclusions

Lancaster University

- Discussed the limitations of the current Intent-based solutions
- Proposed Intent-Based NBI framework
- Proposed a Service Consumer Intent declarative expression along with its corresponding prescriptive policies
- Demonstrated a caching workload intent, and its corresponding policies and refinement in a CCDN use case

Future Work

Lancaster University

- Extend intent-to-policy mapping to be dynamic based on several criteria and map them to existing Microservices.
- Investigate different intent targets.
- Implement the intent-based framework in a real cloud-based testbed along with the required APIs and translations.

Lancaster University

Backup Slides

Different Intent Types

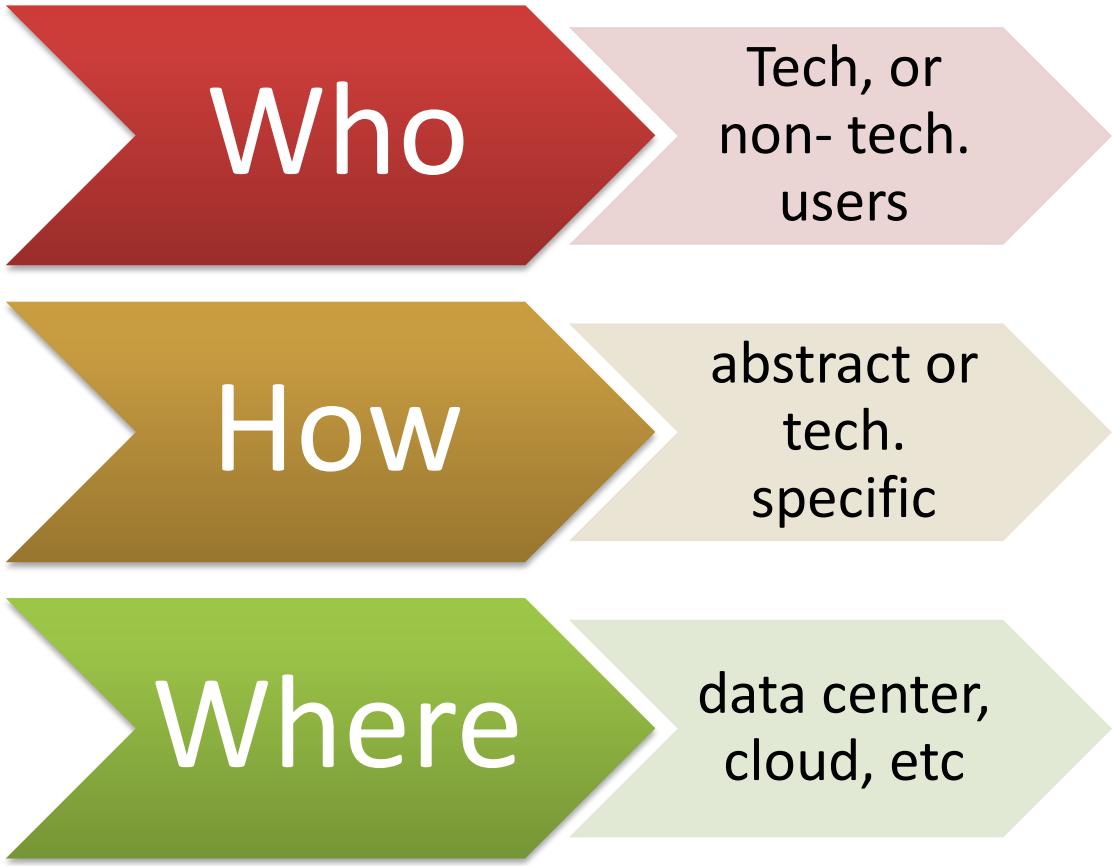

Meta-Analysis

Lancaster University

TABLE I: Summary of the results of our meta-analysis of different intent-based solutions.

Intent-Based Solution	Intent Expression	Domain	Level
Boulder [8]	*Subject, Predicate, Object: {Constraints, Conditions}*	Networking	Presc.
ONOS Intent Framework [9]	*Network Resource, Constraints, Criteria, Instructions*	Networking	Presc.
(NEMO) by Huawei [10]	*Object + Operation* or *Object + Result* (under test and not used yet)	Networking	Presc. Decl.
Group-based Policy (GBP) [11]	*Endpoint group, contract {subject: {rules: {classifier and action set}}}*	Networking / NFV	Presc.
(NIC) by HP [12]	*Source Composite Endpoint, Destination Composite Endpoint, Traffic operation and constraints*	Networking / NFV	Presc.
(DOVE) by IBM [13]	Not specified	Networking / NFV	–
Intent-based virtualisation Platform [14]	*Resources, Conditions, Priority, and Instructions*	Networking / NFV	Presc.
(INSpIRE) [15]	*Traffic Type, Source, Destination, Context level, Contexts list*	Networking / NFV	Presc.
Intent-based NBI service-oriented architecture [16]	application-specific language	Networking	–
(iNDIRA) [17]	*Subject (Service or Condition), Relationship (has Arguments), Objects (multiple parameters)*	Networking	Presc.
(SENSE) [18]	*Service type, Service alias, Connections: {name,terminals, bandwidth: { qos_class, capacity, unit}}, schedule: {start, end, duration}*	Networking / NFV	Presc.
Interactive Intent-based Negotiation Scheme [19]	*Verbs, Nouns, Modifiers*	Networking / NFV	Presc.
(MD-IDN) [20]	*Action, Endpoint 1, Traffic type, Endpoint 2*	Networking	Presc.
Janus system [21]	*Endpoint-Group1, Connection attributes: {protocol, port, bandwidth, latency, middle-box }, Endpoint-Group 2*	Networking	Presc.
Northbound Interface [22]	*Predicate, Commodity, Target (resources), Constraint, Condition*	Networking	Presc.
Adaptive Service Deployment [23]	*Verb, Object, Modifiers, Subject*	General use cases, e.g. storage, caching, IDS	Presc.

Cloud CDN Use Case: ***Caching*** Intent

When an intent developer wants to create a new intent, he has to decide how users can express their high-level targets using the Service-oriented declarative intent expression.

<SERVICE> <RESOURCES> <CONJUNCTION> <TARGET>

For example if **<TARGET>** is **<WORKLOAD>**

The intent developer has to determine how this could be expressed by the CP

either *numerically* such as ***"I want Caching for Content x to meet 10,000 requests/region"***

or describe it with an *adjective* such as ***"I want Caching for Content x with the maximum requests/region"***.

Caching Intent Expression Descriptors

Lancaster University

TABLE II. BASIC EXPRESSION OF CACHING SERVICE-ORIENTED INTENT <SERVICE> <RESOURCES> <CONJUNCTION>

Tag	*Basic Expression*
<SERVICE>	Caching
<RESOURCES>	contents to be cached
<CONJUNCTION>	that can handle / that can meet / etc.

TABLE III. BASIC EXPRESSION SYNTAX OF CACHING SERVICE-ORIENTED INTENT **<TARGET>** AS **<WORKLOAD>**

Tag	*Basic Expression*
<WORKLOAD>	**<NUMBER>** **<UNIT>** or **<ADJECTIVE>** **<UNIT>** or **<ADJECTIVE>** **"workload"**

TABLE IV. BASIC EXPRESSION **<WORKLOAD>**

Tag	*Basic Expression*
<NUMBER>	numeric values that can represent the workload
<UNIT>	requests/region / requests/sec / etc.
<ADJECTIVE>	max / min / dynamic / high / medium / etc.

Cloud CDN Use Case: ***Caching*** Intent

Lancaster University

Then, the intent developer determines how to decompose the ***declarative intent*** expressed by the CP to a set of ***abstract policies***

If a ***new caching request*** is received
    then ***allocate cache servers*** with ***maximum storage capacity***
< sequential >
If ***new user request*** is received
    then ***forward request*** to ***cache server*** with ***least latency***

TABLE V. BASIC EXPRESSION OF **<POLICY>** SYNTAX FOR **<WORKLOAD>** TARGET

Tag	*Basic Expression*
<CONDITIONS>	new caching request /incoming user request /etc.
<ACTIONS>	allocate cache servers() /forward request() /etc.
<CONSTRAINTS>	with max storage / with least latency / etc.
<OPERATOR>	parallel / sequential

Intent Refinement

	Abstract Policies

Initial Intent's Policies	**Policy 1:**
	Condition: if a ***new caching request*** has been received
	Action: proactively allocate caches
	Constraints: ***number*** of caches that can handle the average requested load
	Operator: **Sequential**
	{
	Policy 2:
	Condition: if the caches' ***thresholds*** have been exceeded
	Action: scale out the cluster
	Constraints: ***number*** of caches to be started
	Operator: **Parallel**
	Policy 3:
	Condition: if the caches are ***underutilized***
	Action: scale in the cluster
	Constraints: ***number*** of caches to be stopped
	}

Refined Intent's Policies	{
	Policy 1:
	Condition: if the current day is a ***weekday*** AND the caches' optimistic ***thresholds*** have been exceeded
	Action: scale out the cluster
	Constraints: ***at time*** X_i
	Operator: **Parallel**
	Policy 2:
	Condition: if the current day is a ***weekday*** AND the caches are ***underutilized***
	Action: scale in the cluster
	Constraints: ***underutilization time*** $> X_j$
	}
	Operator: **Parallel**
	{
	Policy 3:
	Condition: if the current day is a ***weekend*** AND the caches' pessimistic ***thresholds*** have been exceeded
	Action: scale out the cluster
	Constraints: ***at time*** X_k
	Operator: **Parallel**
	Policy 4:
	Condition: if the current day is a ***weekend*** AND the caches are ***underutilized***
	Action: scale in the cluster
	Constraints: ***underutilization time*** $> X_m$
	}

	Technical Policies
Initial Intent's Policies	**Policy 1:** **Condition:** if a ***new caching request*** has been received **Action:** allocate ***VMs*** as caches **Constraints:** ***1 VM*** **Operator:** **Sequential** **Policy 2:** **Condition:** if the caches' ***CPU Utilization > 80%*** **Action:** scale out the cluster by spinning up VMs **Constraints:** ***1 VM*** **Operator:** **Parallel** **Policy 3:** **Condition:** if the caches' **CPU Utilization < 20%** **Action:** scale in the cluster by stopping VMs **Constraints:** ***1 VM***

Refined Intent's Policies	{
	Policy 1:
	Condition: if the current day is a ***weekday*** AND the caches' ***CPU Utilization > 75%*** have been exceeded
	Action: scale out the cluster
	Constraints: ***at 14:45***
	Operator: **Parallel**
	Policy 2:
	Condition: if the current day is a ***weekday*** AND the caches' **CPU Utilization < 20%**
	Action: scale in the cluster
	Constraints: ***underutilization time > 20 minutes***
	}
	Operator: **Parallel**
	{
	Policy 3:
	Condition: if the current day is a ***weekend*** AND the caches' ***CPU Utilization > 70%***have been exceeded
	Action: scale out the cluster
	Constraints: ***at 14:35***
	Operator: **Parallel**
	Policy 4:
	Condition: if the current day is a ***weekend*** AND the caches' **CPU Utilization < 20%**
	Action: scale in the cluster
	Constraints: ***underutilization time > 30 minutes***
	}

Lancaster University

Thank You

Shiyam Alalmaei, Matthew Broadbent, Yehia Elkhatib, Mehdi Bezahaf, and Nicholas Race

School of Computing and Communications, Lancaster University, United Kingdom

Email: {s.alalmaei, m.broadbent, y.elkhatib, mehdi.bezahaf, n.race}@lancaster.ac.uk