IoT Analytics at the Edge

Alan Clark, CEO, Telchemy
alan.d.clark@telchemy.com

IoT Analytics – useful data or data mountain?

- Provides a useful data for:
  - Device control
  - Usage statistics
  - Marketing and demographics
- Some applications need high resolution data
  - frequent reporting creates massive load
- Analytics reporting should provide detailed and accurate data – comes with consequences
  - Database/disk storage issues
  - Network load

telchemy
Actively managing multimedia

IoT analytics - economics

Number of "things"	Reporting Interval	Database IOPS	Database Gbytes/24 hrs
10,000	15 minutes	11	2
10,000	30 seconds	333	43
100,000	15 minutes	111	14
100,000	30 seconds	3,333	432
1,000,000	15 minutes	1,111	144
1,000,000	30 seconds	33,333	4,320
10,000,000	15 minutes	11,111	1,440
10,000,000	30 seconds	333,333	43,200

Assumes single 1500 byte IP packet per report

telchemy
Actively managing multimedia

Datanomics

- We store
  - 2,500,000 Terabytes per day
  - Which is 300 Mbytes per person per day
- What gets stored?
  - 1 billion Credit/Debit card transactions - 10 Tb?
  - 10 billion phone call records – 20 Tb?
  - ……..
- Does it get used?
  - There are 10-20 million marketing people on the planet, they would each need to analyze 100 Gbytes of data per day

telchemy
Actively managing multimedia

IoT Analytics models

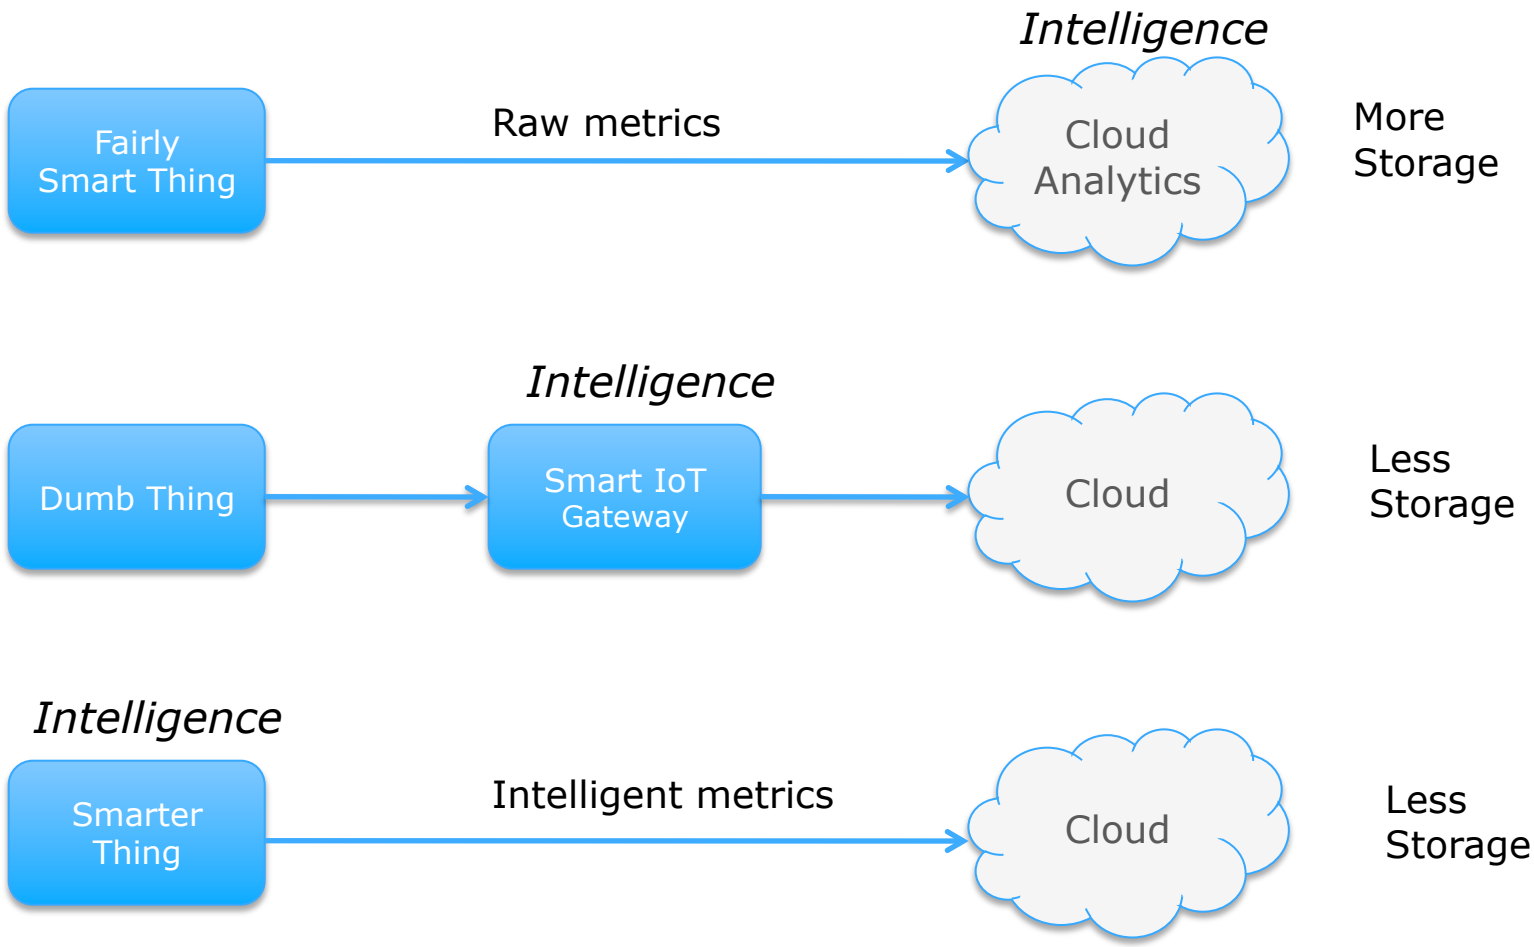

telchemy
Actively managing multimedia

IoT Analytics agents

- Custom agent
  - Historically most agents were simply “part of the code” and very simple
  - Metrics vary by vendor and by application
  - Some agents are more sophisticated – e.g. AMI
- “Connectors”
  - Typically very simple agents that allow locally available metrics to be sent to management system
- Third party agent
  - Licensable 3rd party agent for specific application
  - Examples VQmon VoIP agent, RMON network agent
  - Enables “smarter data”

telchemy
Actively managing multimedia

IoT Analytics Agent

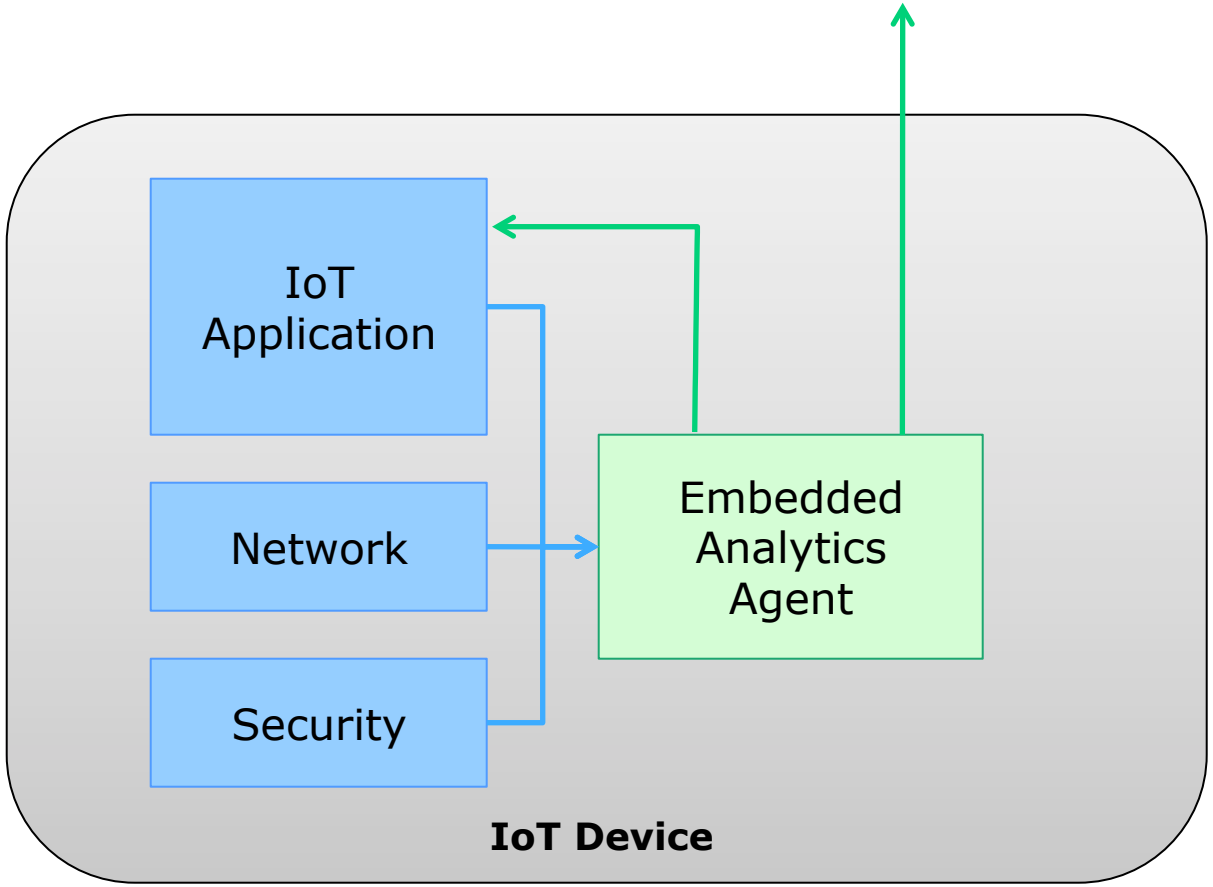

telchemy
Actively managing multimedia

IoT Analytics at scale

- Sample raw data frequently to get resolution needed to understand application
- Summarise raw data, send less frequent reports
- Locally correlate/ analyze raw metrics and send intelligent metrics
- Use exception based reporting and probabilistic alerting to provide fast indication of problems

telchemy
Actively managing multimedia

Use case 1 – electric meter

Smart Electric
Meter with AMI

AMR - monthly data collection for billing

AMRplus – hourly interval data – daily collection
- improved usage information

AMI – 5 minute interval data - hourly collection
- usage information, remote problem “triage”
- “last gasp” outage detection

telchemy
Actively managing multimedia

Use case 1 – electric meter

99.975% of data is not related to outages

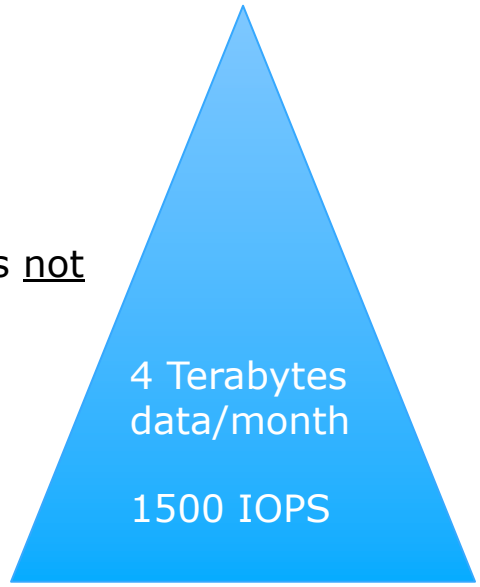

Example:

2.5 million customer utility using AMI

Supports billing, maintenance

Enables customers to review usage

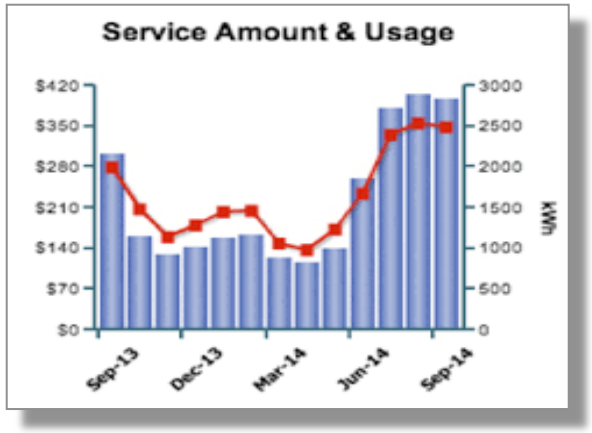

telchemy
Actively managing multimedia

Use case 1 – electric meter

1Gbyte data/month relates to outages

99.975% of data is <u>not</u> related to outages

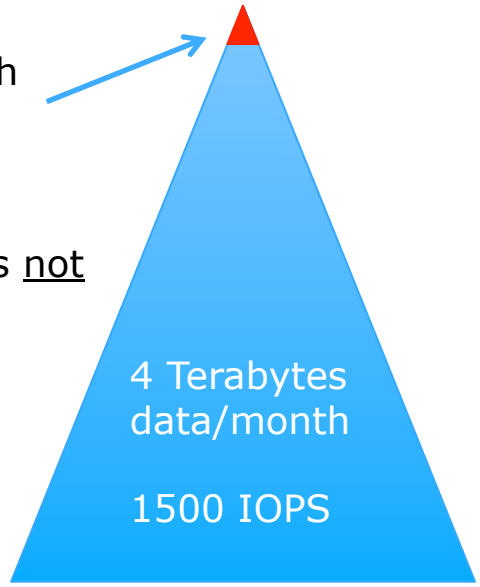

Duke University suggesting 6 second intervals - would increase this to 200 Terabytes/month and 50,000 IOPS

Example:

2.5 million customer utility using AMI

Supports billing, maintenance

Enables customers to review usage

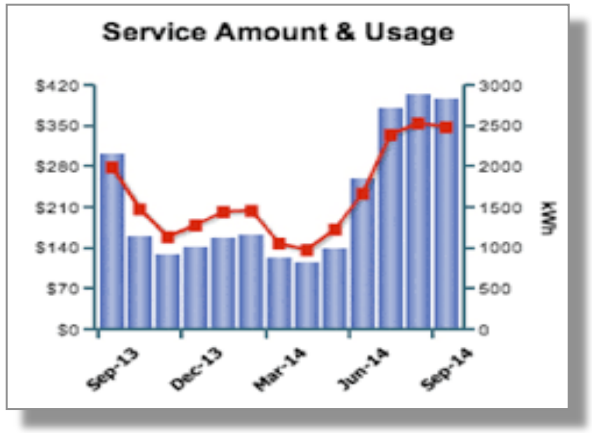

telchemy
Actively managing multimedia

Use case 2 – VoIP endpoint

IP Phones

Residential Gateways

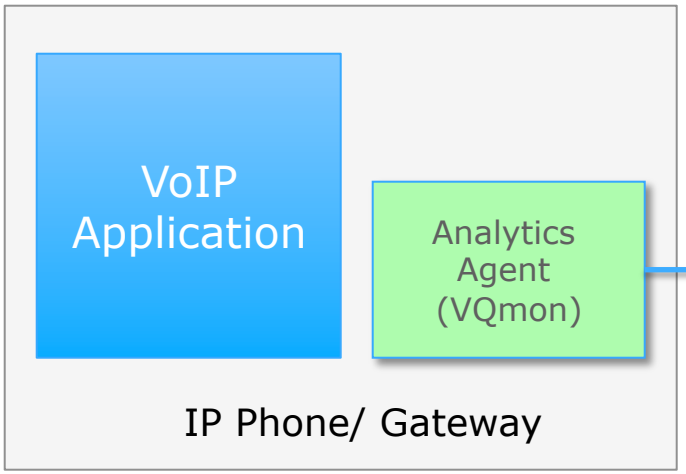

RFC 6035

QoE report sent at end of call through SIP infrastructure

Estimated 500 million+ VoIP endpoint agents in use

telchemy
Actively managing multimedia

Use case 2 – VoIP endpoint

- Problem – impairments are time varying and sending average loss, jitter… does not allow accurate QoE measurement

- Naïve solution (*actually used by one vendor !!!*)
  - Send report every 5 seconds for every active call
  - 40 reports per call – 40x network bandwidth, 40x data storage, 40x IOPS

- Smart solution (300 million copies in use)
  - Use model that learns time distribution of IP impairments (Markov Model)
  - Allows QoE to be calculated accurately at end of call
  - 1 report per call – 1x data storage, 1x IOPS

telchemy
Actively managing multimedia

Use case 2 – VoIP endpoint

Smart Solution

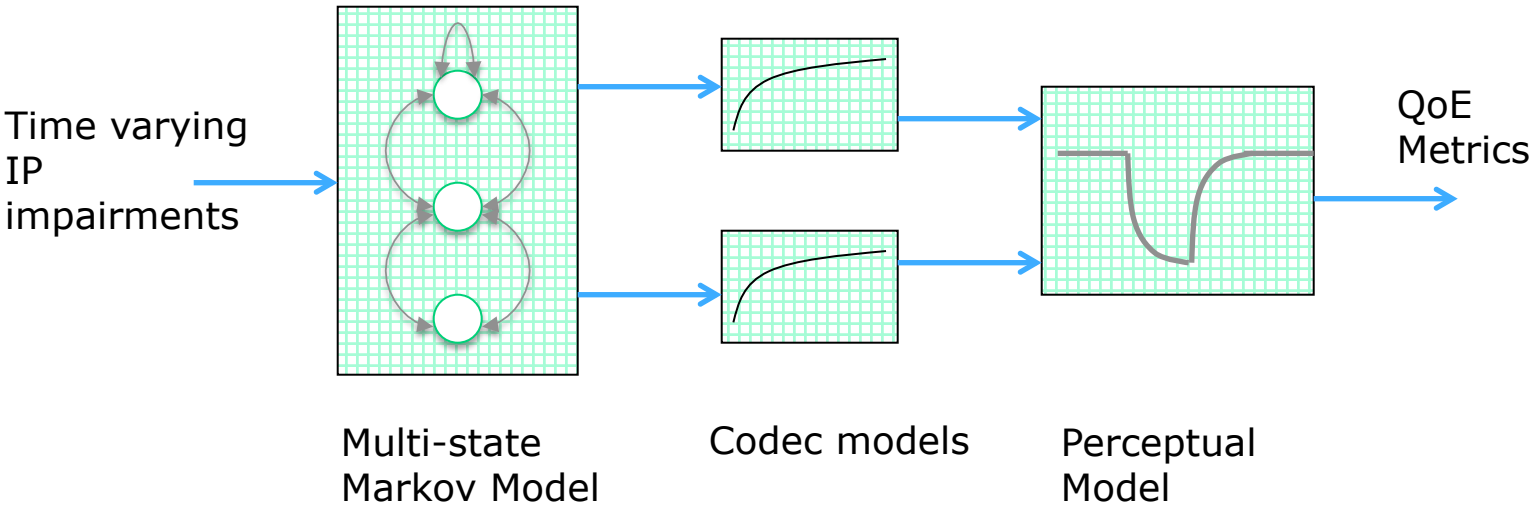

Uses complex algorithms to accurately estimate perceptual quality…
but only uses 0.001 MIPS

Output metrics are accurate and detailed but can be sent in one IP packet

telchemy
Actively managing multimedia

IoT Analytics Agent – ideal characteristics

- Generally:-
  - Small and efficient
  - Accurate and detailed metrics
  - Metrics tailored to device and application
  - Extremely efficent reporting without loss of detail

- Ideally:-
  - Developing an agent should not require in-depth expertize on statistics and reporting
  - Developing an agent should require minimal software development
  - Intelligent metrics play a part in device control

telchemy
Actively managing multimedia

IoT Analytics Agent

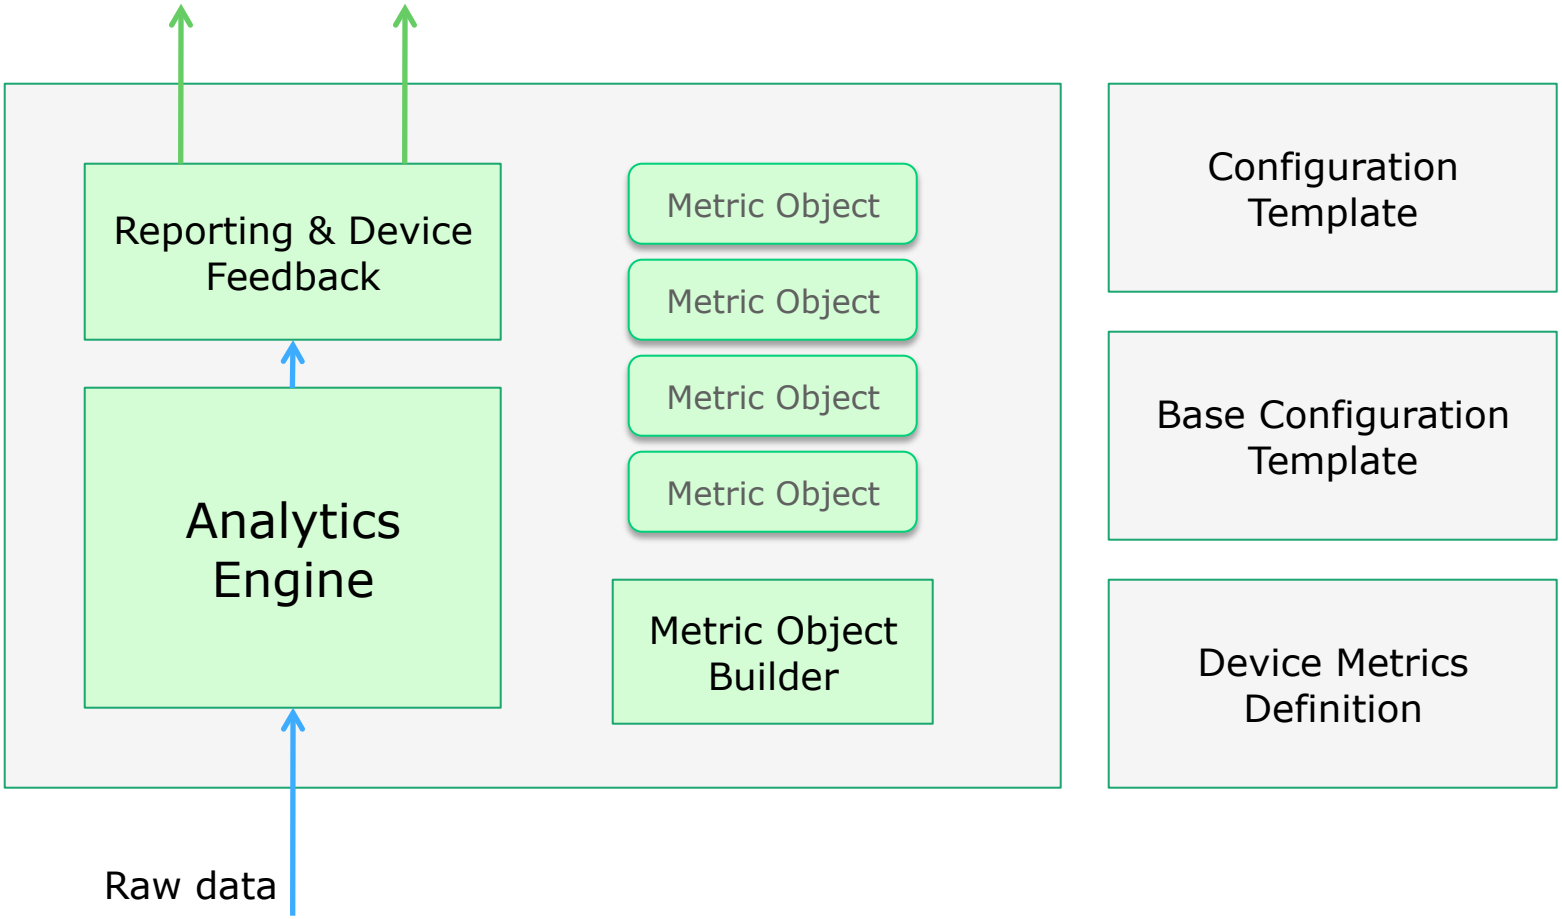

telchemy
Actively managing multimedia

IoT Analytics Agent – compute graph

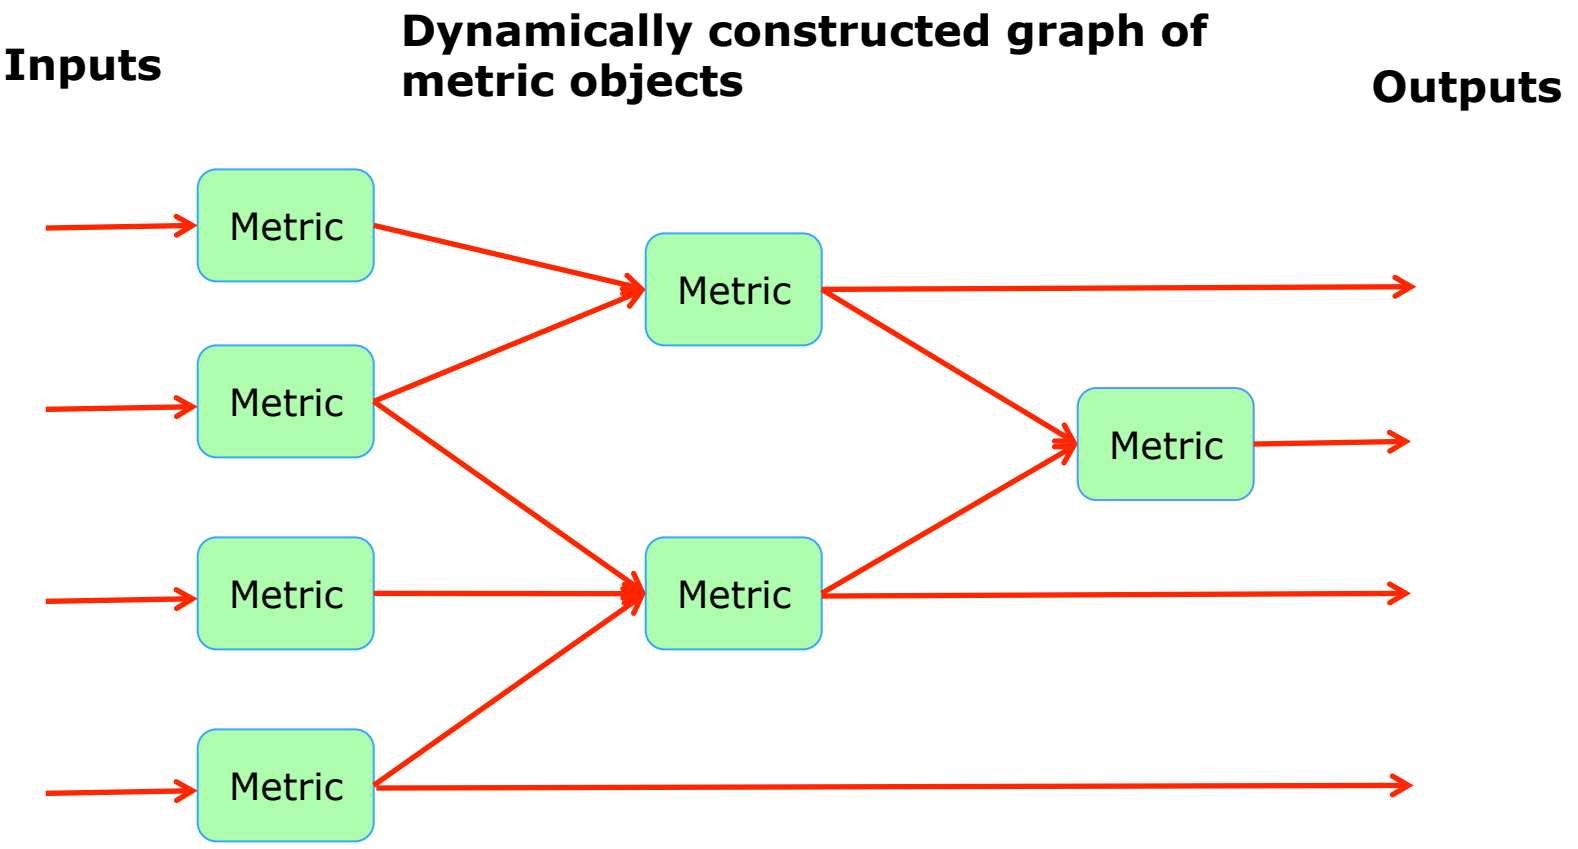

telchemy
Actively managing multimedia

IoT Analytics Agent – Metric Types

- Simple – single output value
  - Unitary – e.g. running average of input
  - Math functions – e.g. log, sqrt…
  - Stats functions – e.g. std dev, trend..
  - Filter functions – e.g. fir, poly
  - Referenced by metric name

- Complex – set of output values
  - Threshold object – time above/ below threshold….
  - Histogram – bin counts, median, mode…
  - Correlation object – bivariate correlation
  - Referenced by metric name "." parameter

telchemy
Actively managing multimedia

IoT Analytics Agent – Compute model

- Embedded real time stream processing model
- Designed for high throughput, very low computational complexity/ code/ memory
- Addresses order of computation, i.e. if an input value is updated then should this trigger an output change
- Addresses missing input data, i.e. if an input value is not updated (sensor failed?) then how should computation be handled?
- Multithreaded – allows input, computation and reporting to be concurrent and asynchronous

telchemy
Actively managing multimedia

IoT Analytics Agent - reporting

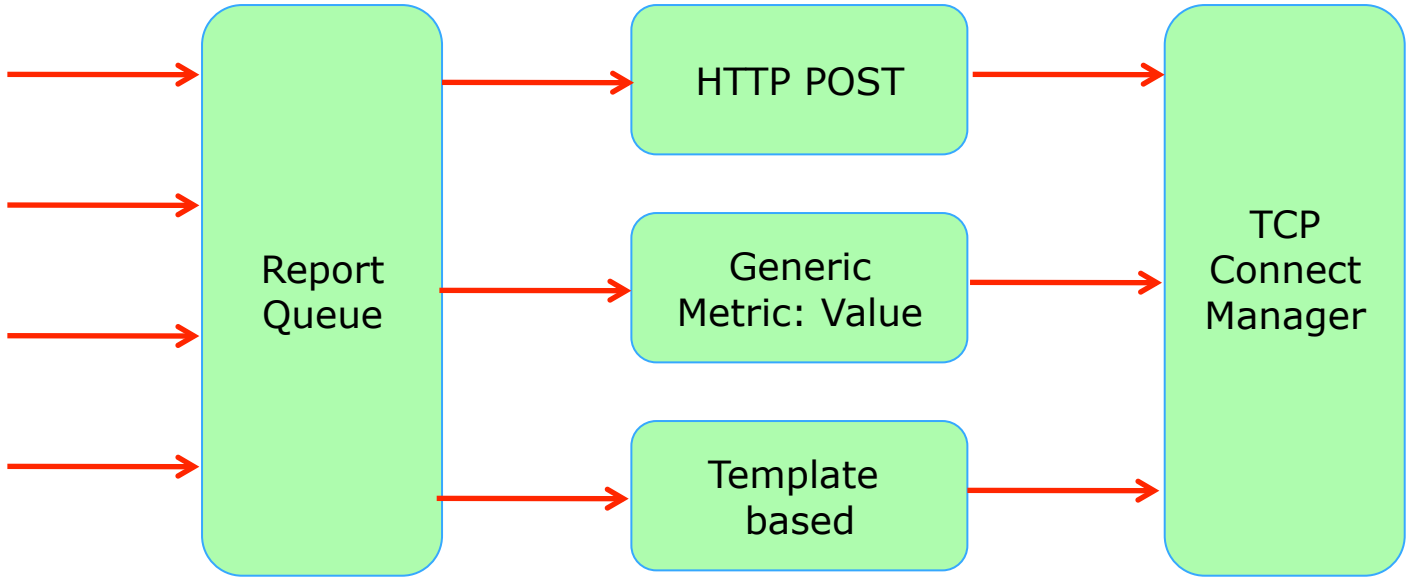

telchemy
Actively managing multimedia

Example configuration file

```
vendor: Telchemy
device_type: mydevice
device_id: 123456
device_version: 1.0

input_port: 8888

report_url: /sandbox/iot/iot_handler.php
report_host: www.resulthostcom
report_port: 80
report_type: POST_URL_encoded
report_interval: 5

metric: temperature, float, input, output
metric: humidity, integer, input, output
metric: average temp, float, output
metric: average humidity, integer, output
metric: th, float, output

rule: average humidity = average8( humidity )
rule: average temp = average16(  temperature  )
rule: th = histogram( temperature,20,0,100 )
```


telchemy
Actively managing multimedia

IoT Analytics Agent

- Small footprint – base agent and dynamically allocated objects
- Transparent use of integer or floating point
- Allows more complex analytics with no/ minimal code development
- Supports device control, reporting, alerting
- Multiple/ flexible reporting options

telchemy
Actively managing multimedia

IoT Analytics Agent vs other approaches?

- IoT Analytics Agent
  - Very fast to implement/ change/ experiment
  - Low complexity embedded application
  - Handles threading, reporting, missing metrics….
  - Tested implementations of math/ functions
- Perl/Python Script
  - Moderate time to implement
  - Greater flexibility in implementing algorithms
  - Need to implement threading, reporting, concurrency
- C/C++/Java implementation
  - Longer time to implement
  - Greater flexibility in implementing algorithms
  - Need to implement threading, reporting, concurrency

telchemy
Actively managing multimedia

IoT Analytics – Edge vs Core Intelligence

- Using “Intelligent Metrics” not “Big Data” leads to more useful analytics with less network bandwidth, storage and IOPS
- Intelligent agent technology allows real time correlation of time varying data at source, results in more accurate analytics
- Generic Agent allows sophisticated analytics to be embedded into devices/ things without writing code
- Cloud based systems still provide cross device correlation, reporting
- Extensible performance management applications support analytics collected from anything, providing sophisticated reporting

telchemy
Actively managing multimedia