




atsu.io


Operational Intelligence for the Digital Factory

(AI for HPC Ops)

The digital factory

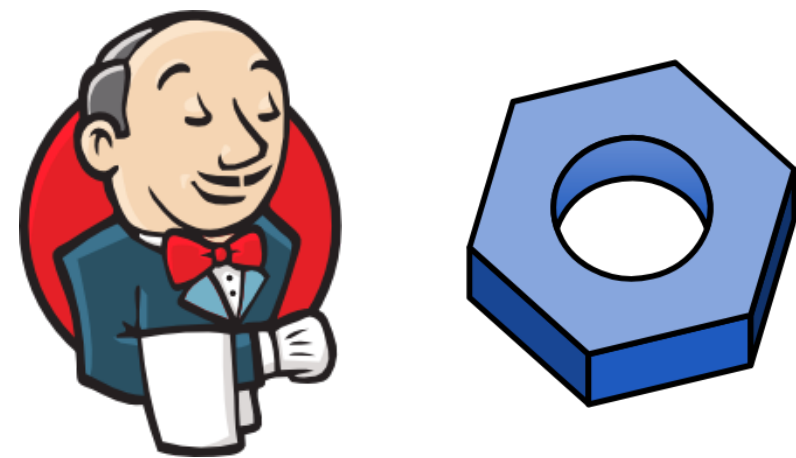


Compute grid




Enterprise storage

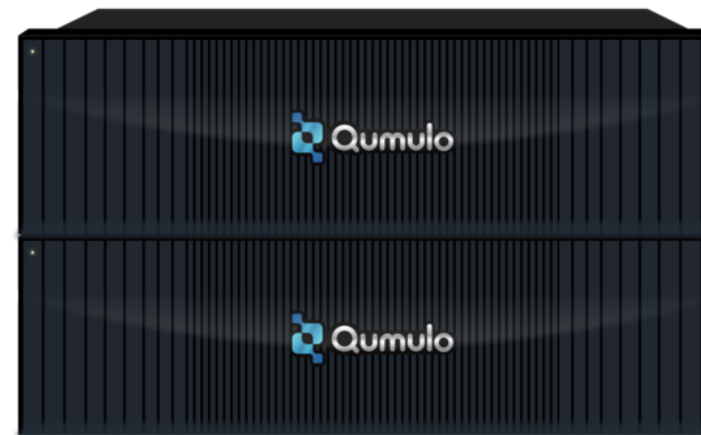
The digital factory



Shared services




Compute grid




Enterprise storage


The digital factory




Shared services



Compute grid



Enterprise storage



quant

ops

pm

dev

Users

Scale of digital production

Hundreds of storage mounts

Thousands of compute hosts


Tens of thousands of user jobs

Millions of application tasks

Billions of I/Os

Tens of billions of core seconds

Quant issues



My home directory is slow.

My workstation is lagging.

My job is taking too long.

Ops issues



Storage is being hammered.

Scheduling database is overwhelmed.

Who is responsible?

PM issues



Will we hit our milestones?

Can we do more of X and less of Y?

Are users properly accounted?

Dev issues



Am I efficiently using shared storage?

Can we pack compute tasks more?

What's the biggest distributed bottleneck?


Factory challenges

Unexpected impacts kill margins.

Idle or misallocated resources sink profits.


Incomplete planning forces expensive mitigations.

Factory goals



Predictable

Reliably submitting, processing, and completing work.



Efficient

Maximizing work completed given available resources.



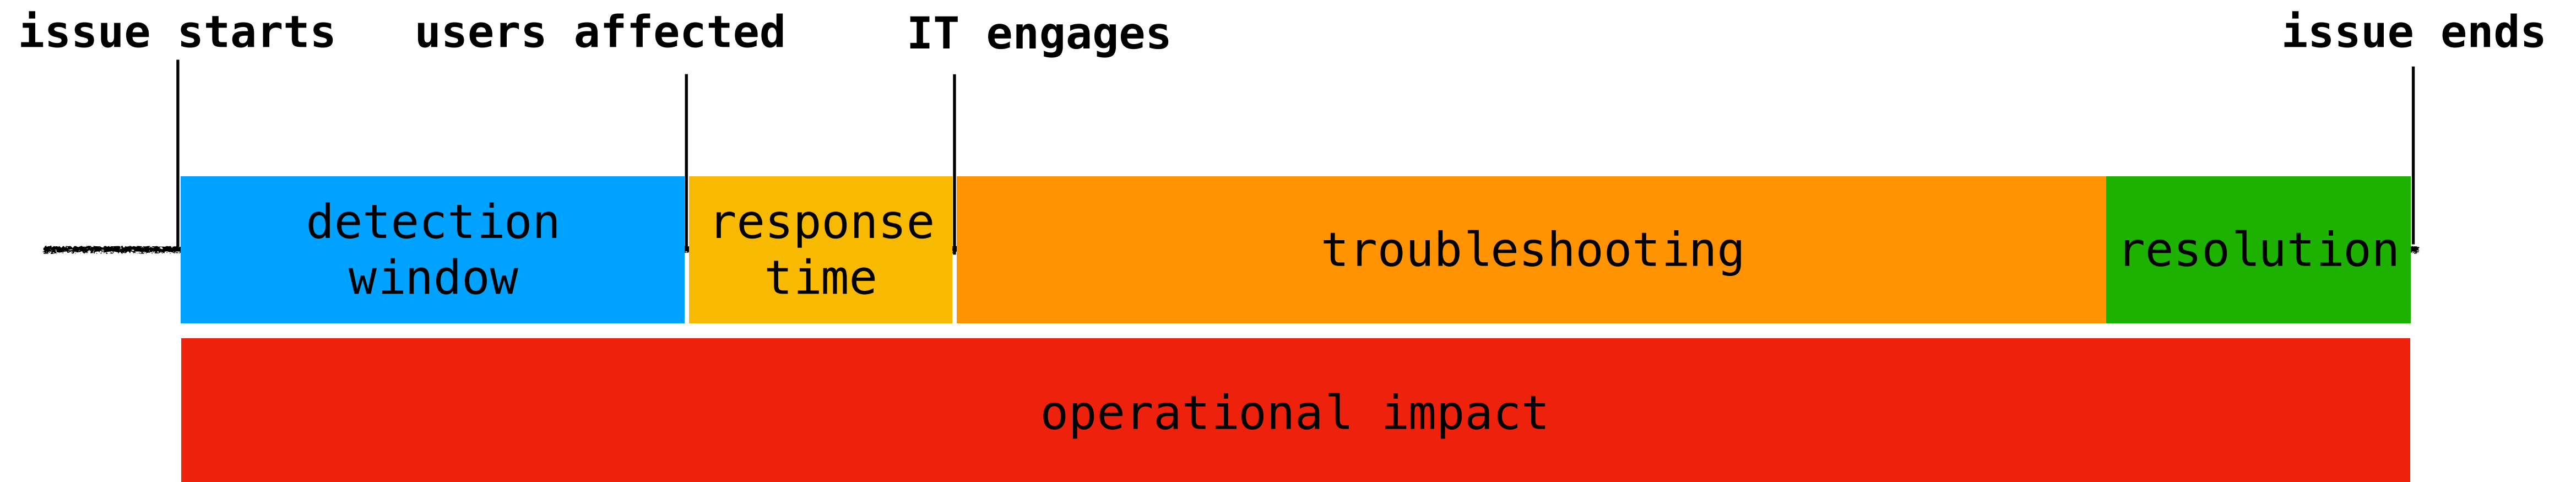
Agile

Understanding the cost/benefit before implementation.

Reducing impact

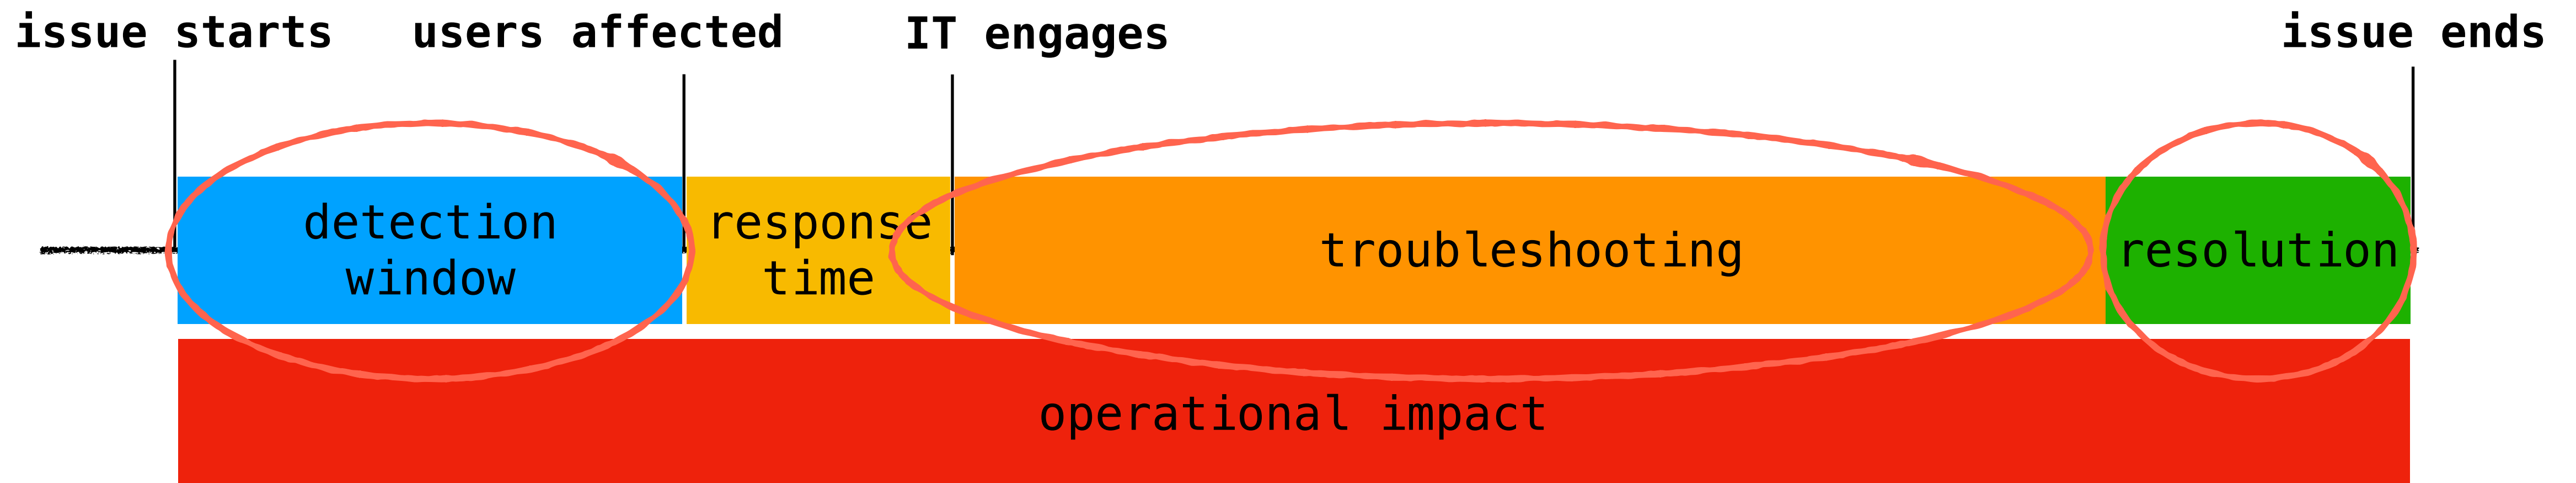


Reducing impact



How do we reduce the operational impact as we scale?

Reducing impact




**Shorten
the
detection
window**

**Provide a list
of causes**

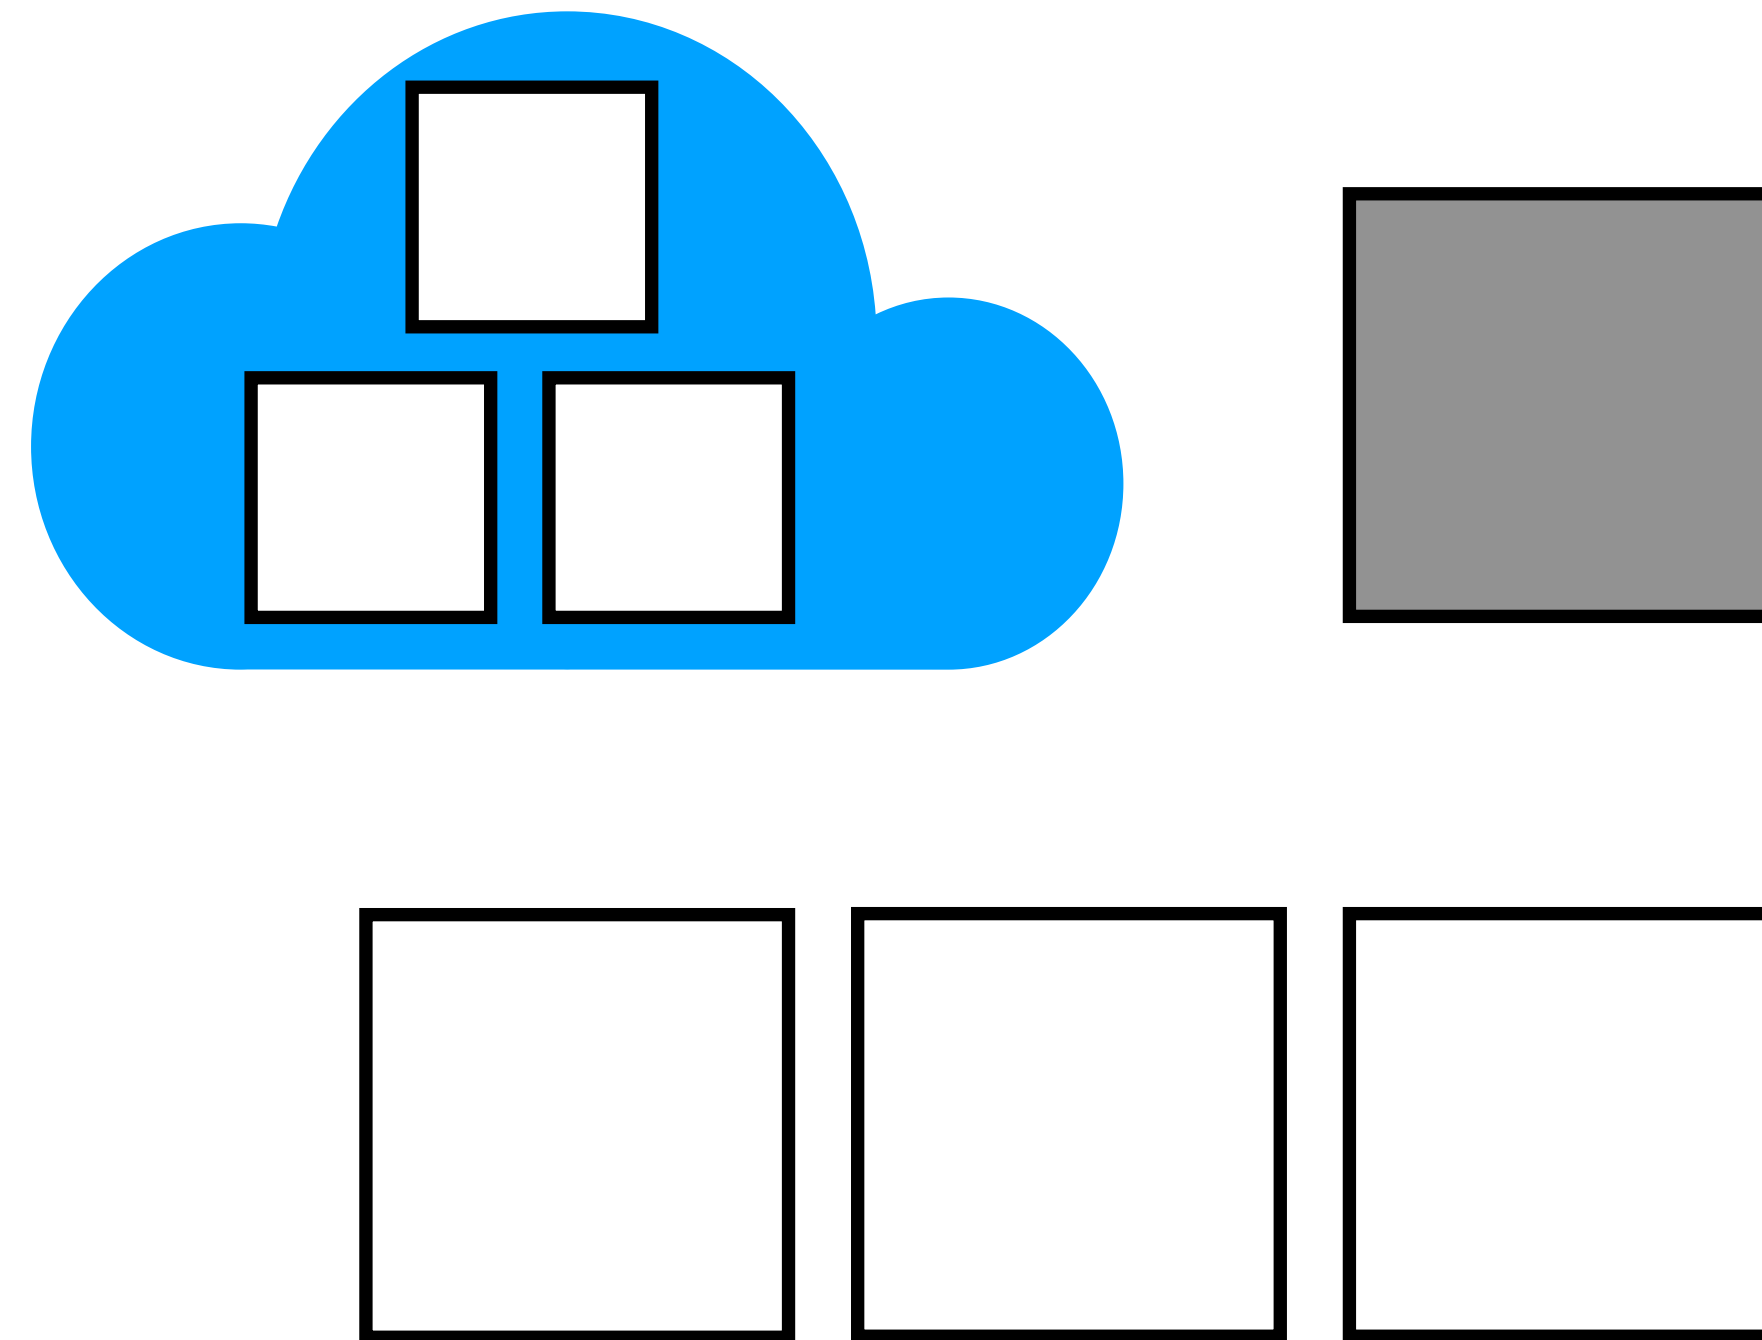
**Automate
resolutions**

Operational Intelligence: training the machines to watch the machines


Improving efficiency



Work




Capacity

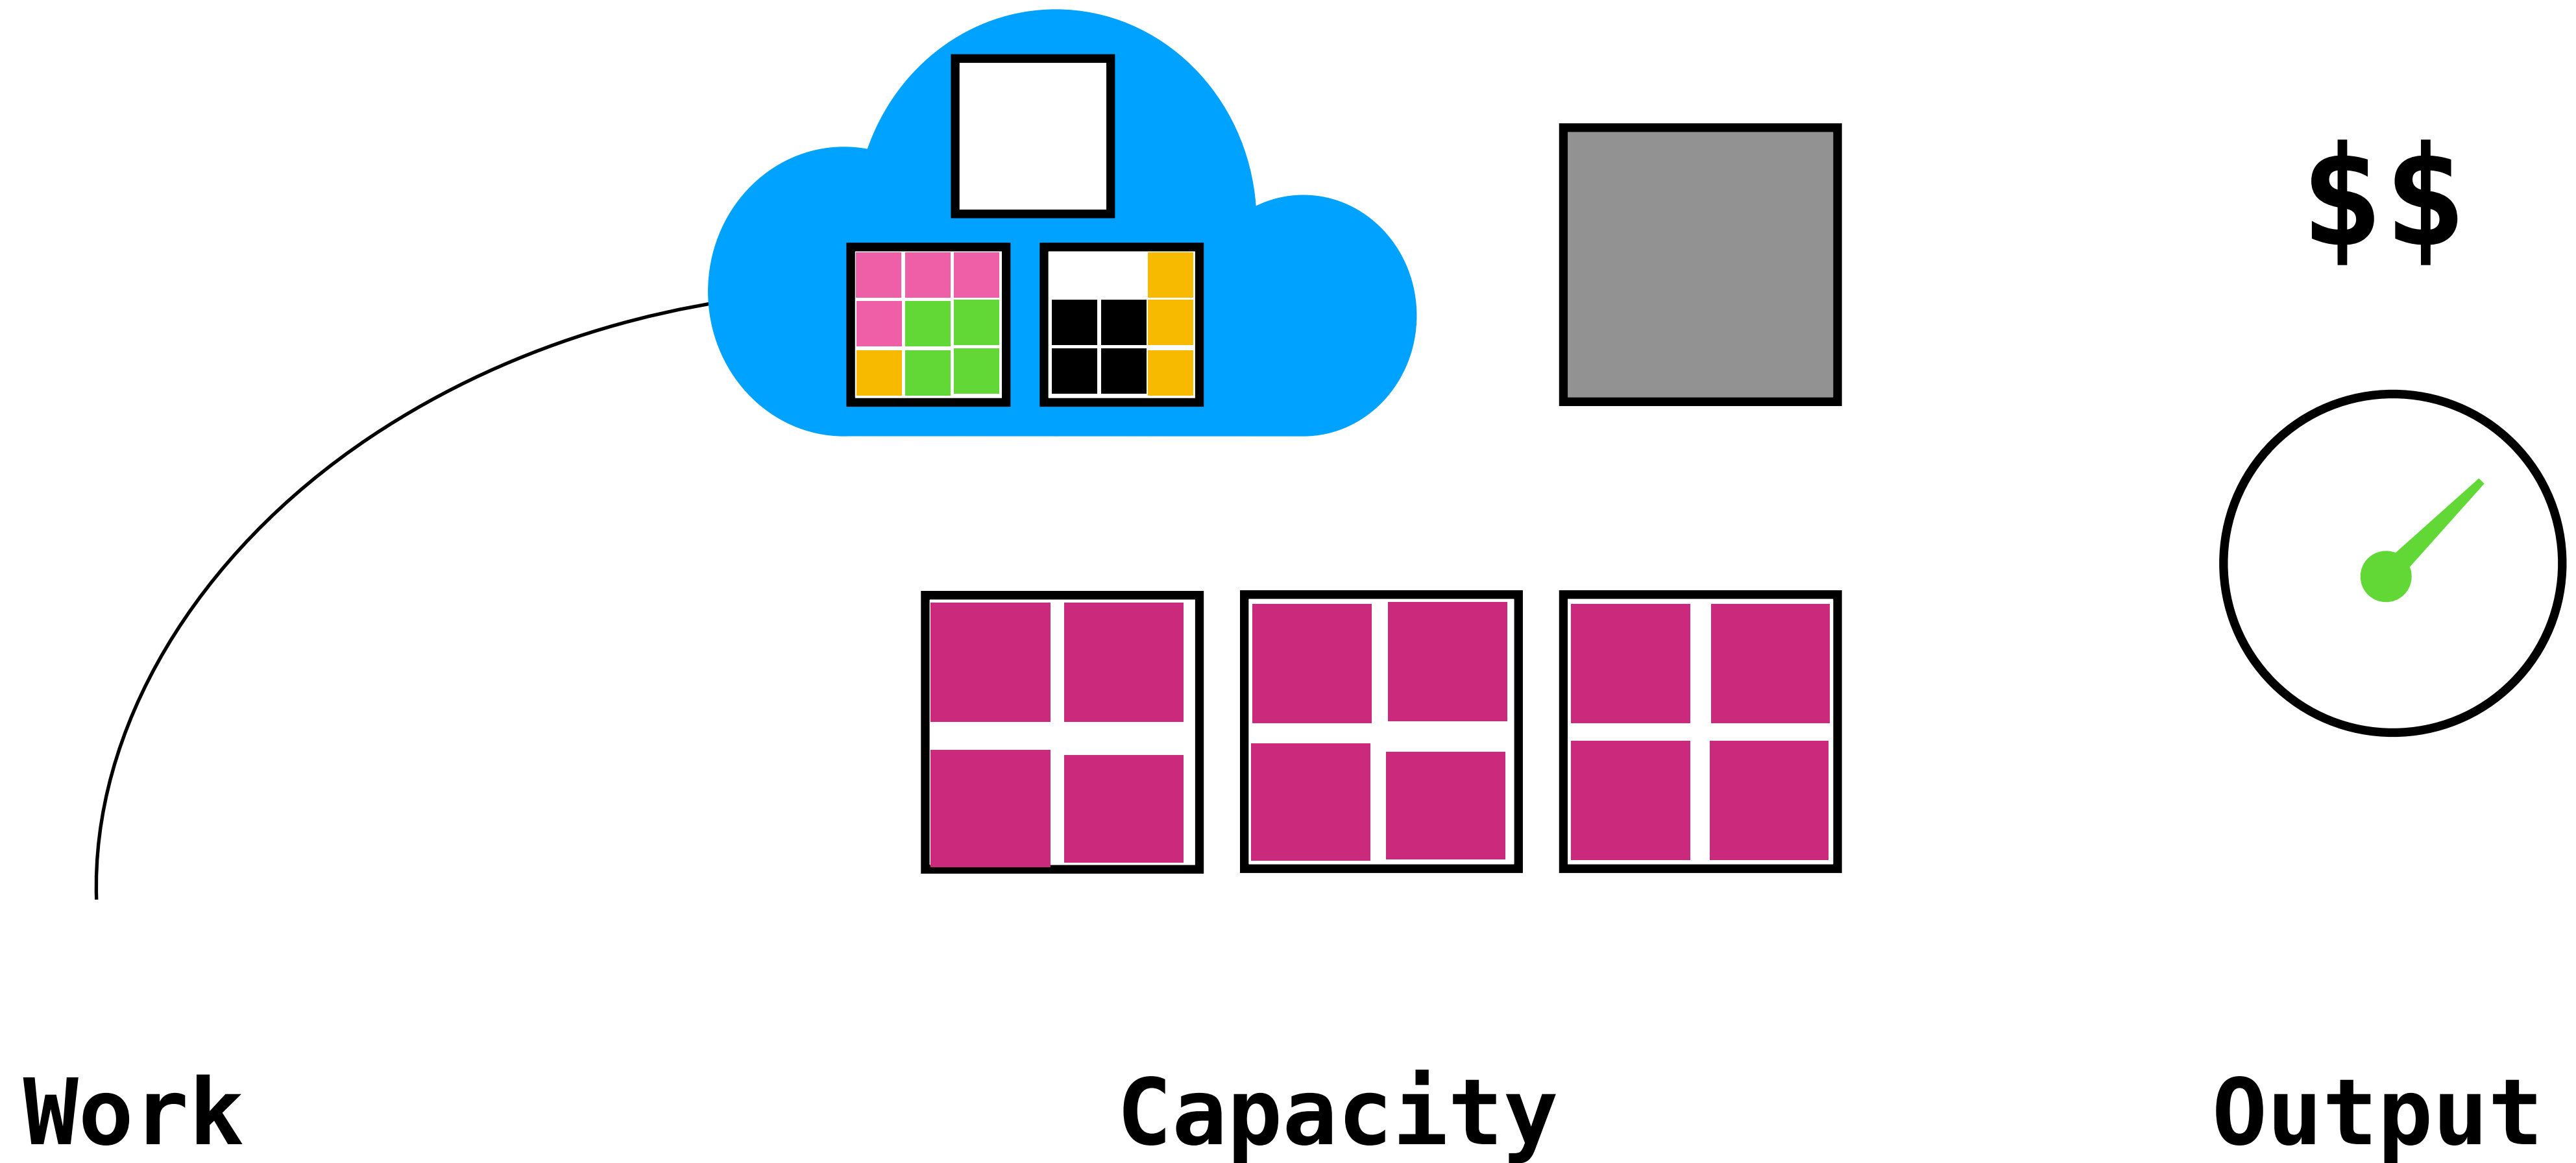


Output

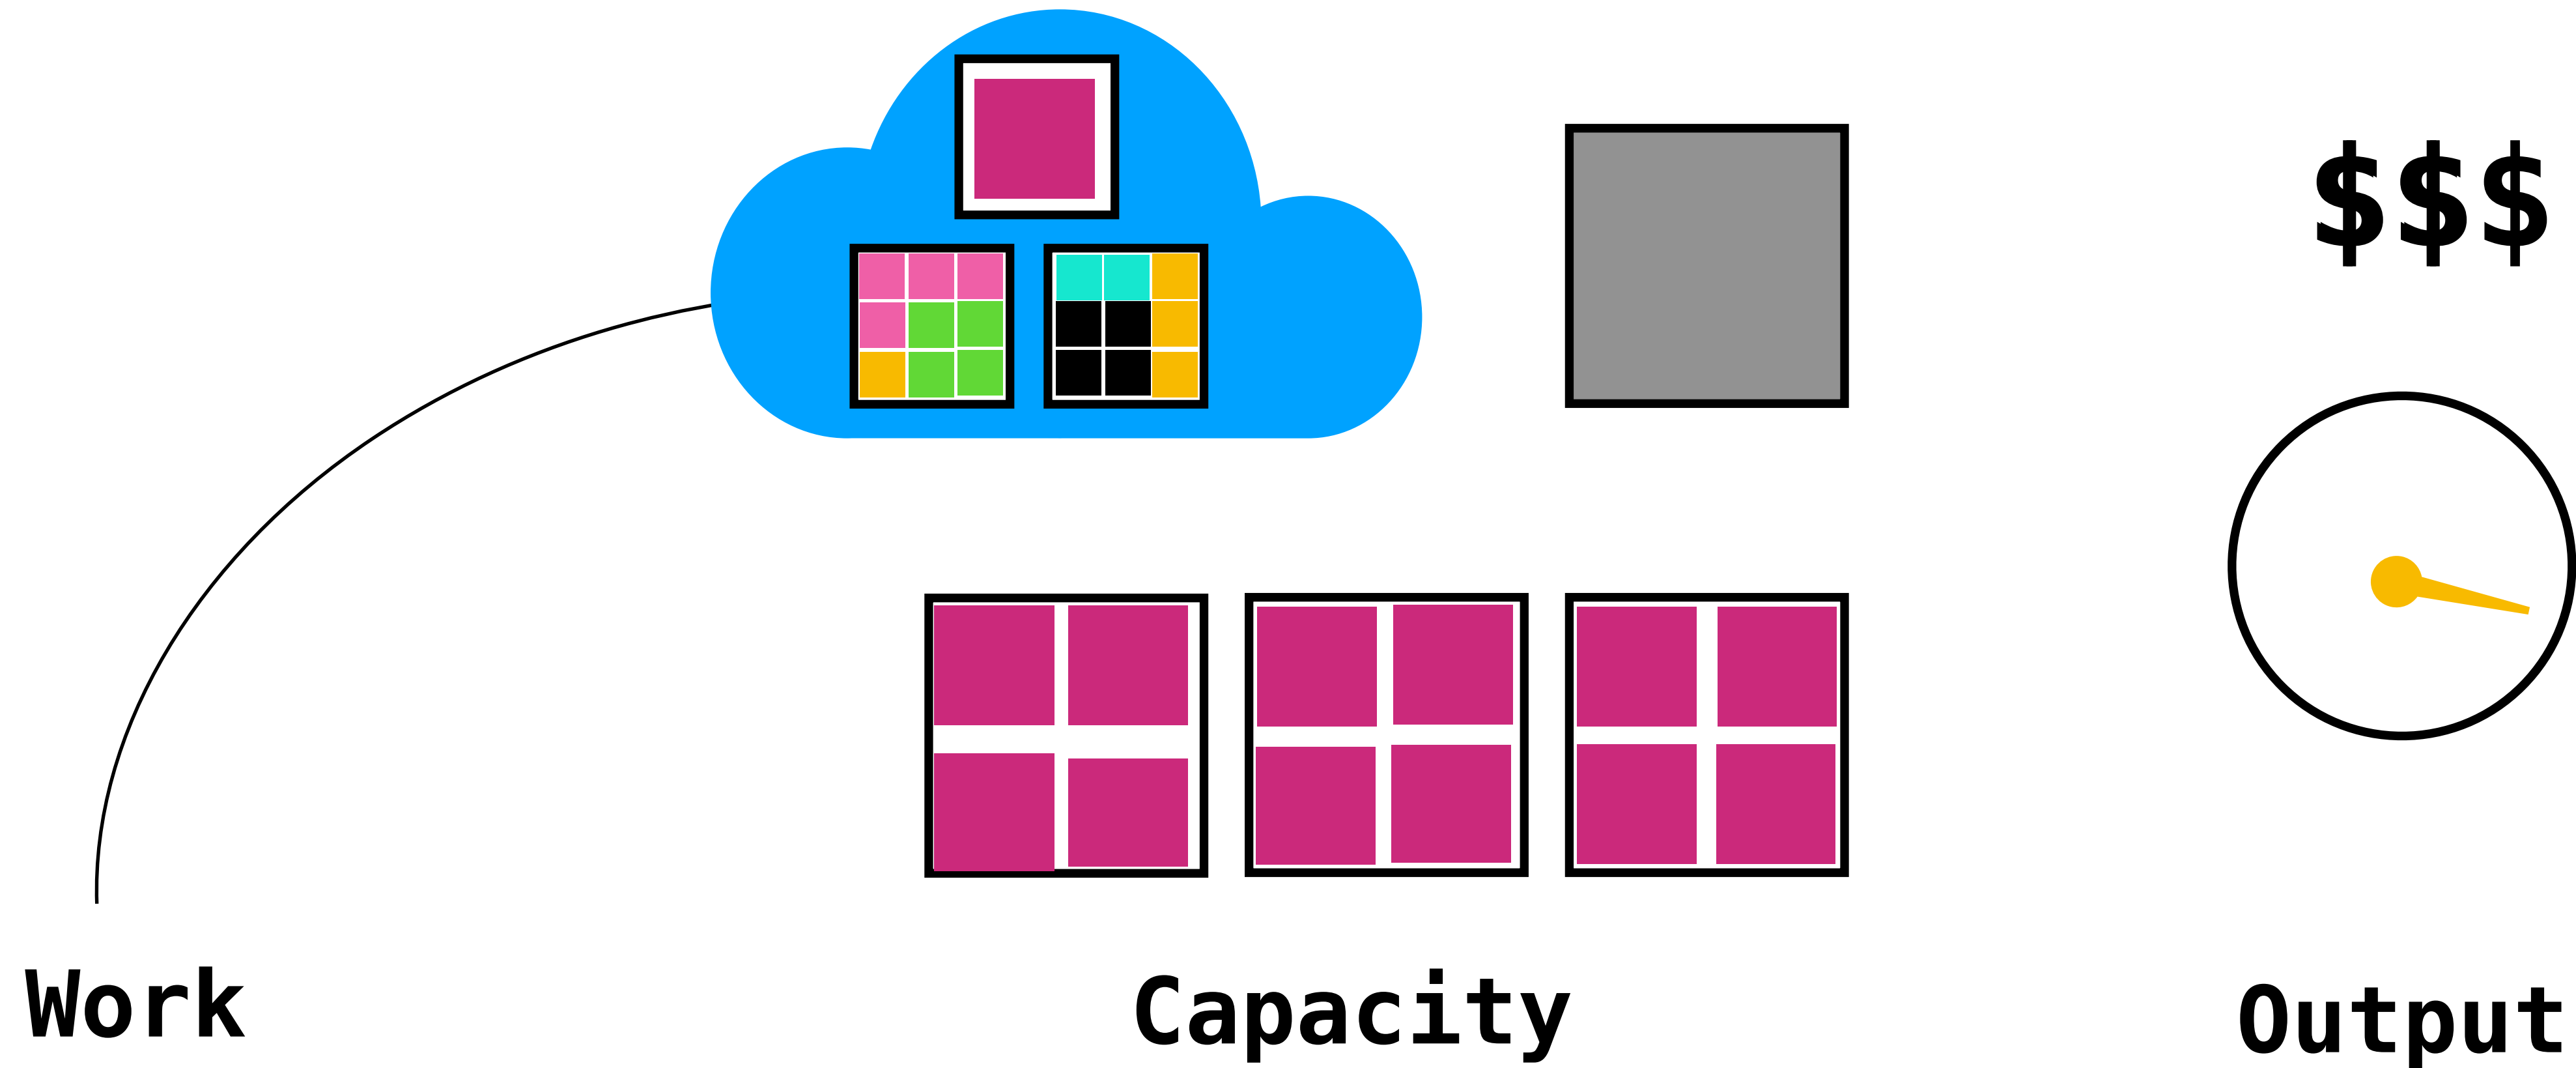
Improving efficiency



Improving efficiency



Improving efficiency



Operational Intelligence: observing the environment leads to better provisioning

Learned agility

- Using historical patterns to extrapolate futures
- Comparing applications and infrastructure over time
- Leveraging trained models to help with repetitive tasks

Operational Intelligence: using your workflow to predict your future workflow

Operational Intelligence


Training the machines to watch the machines

Observing the environment leads to better provisioning

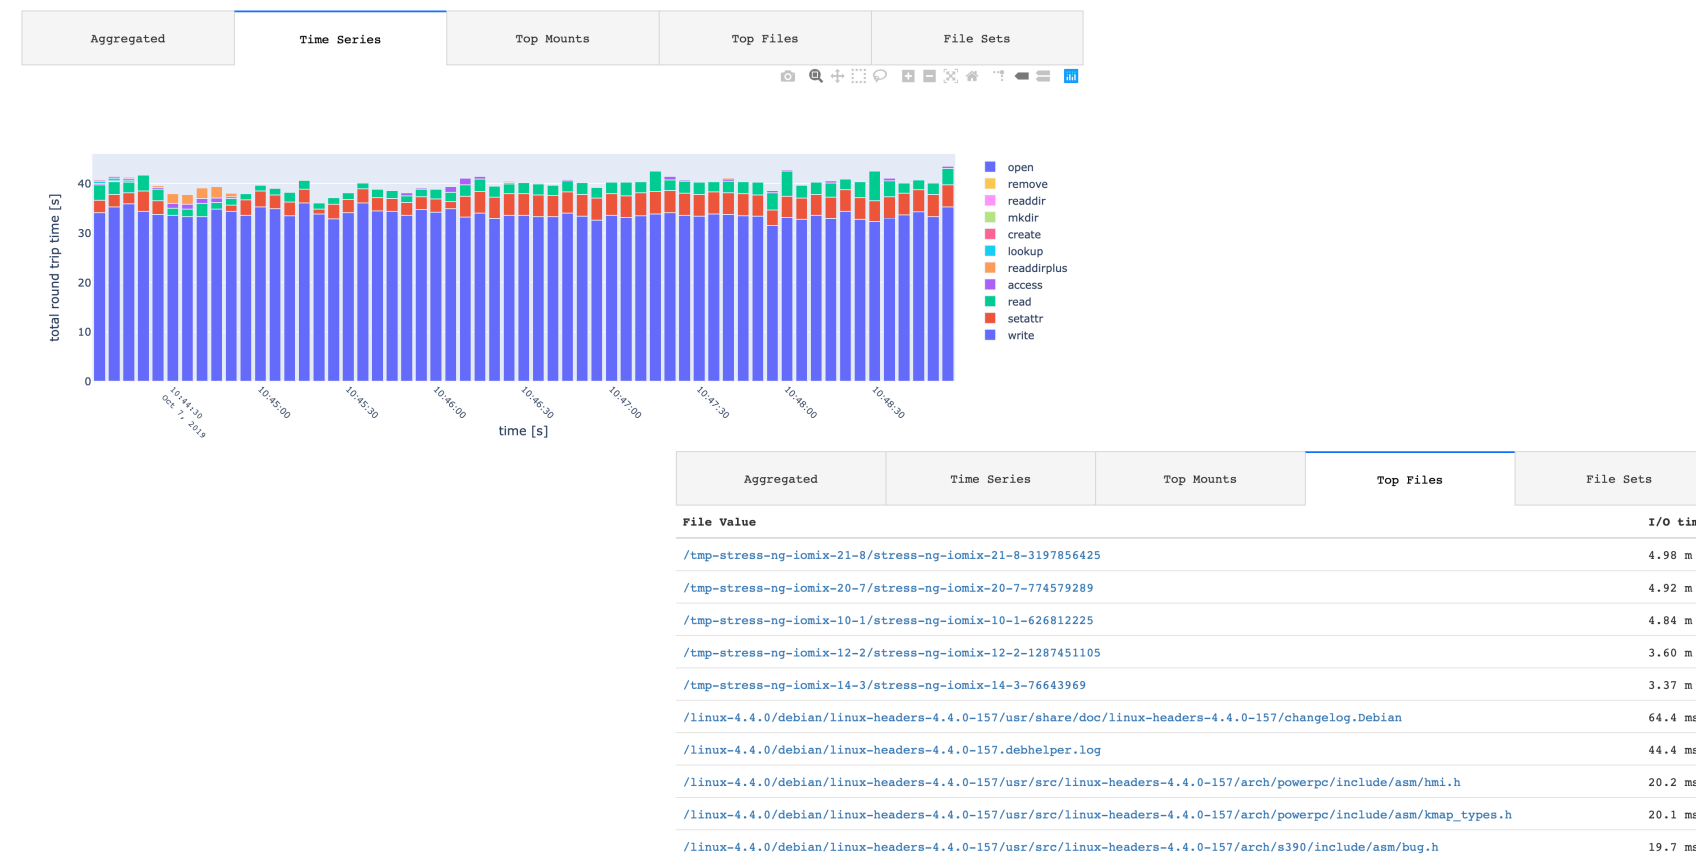
Using your workflow to predict your future workflow

atsu solutions

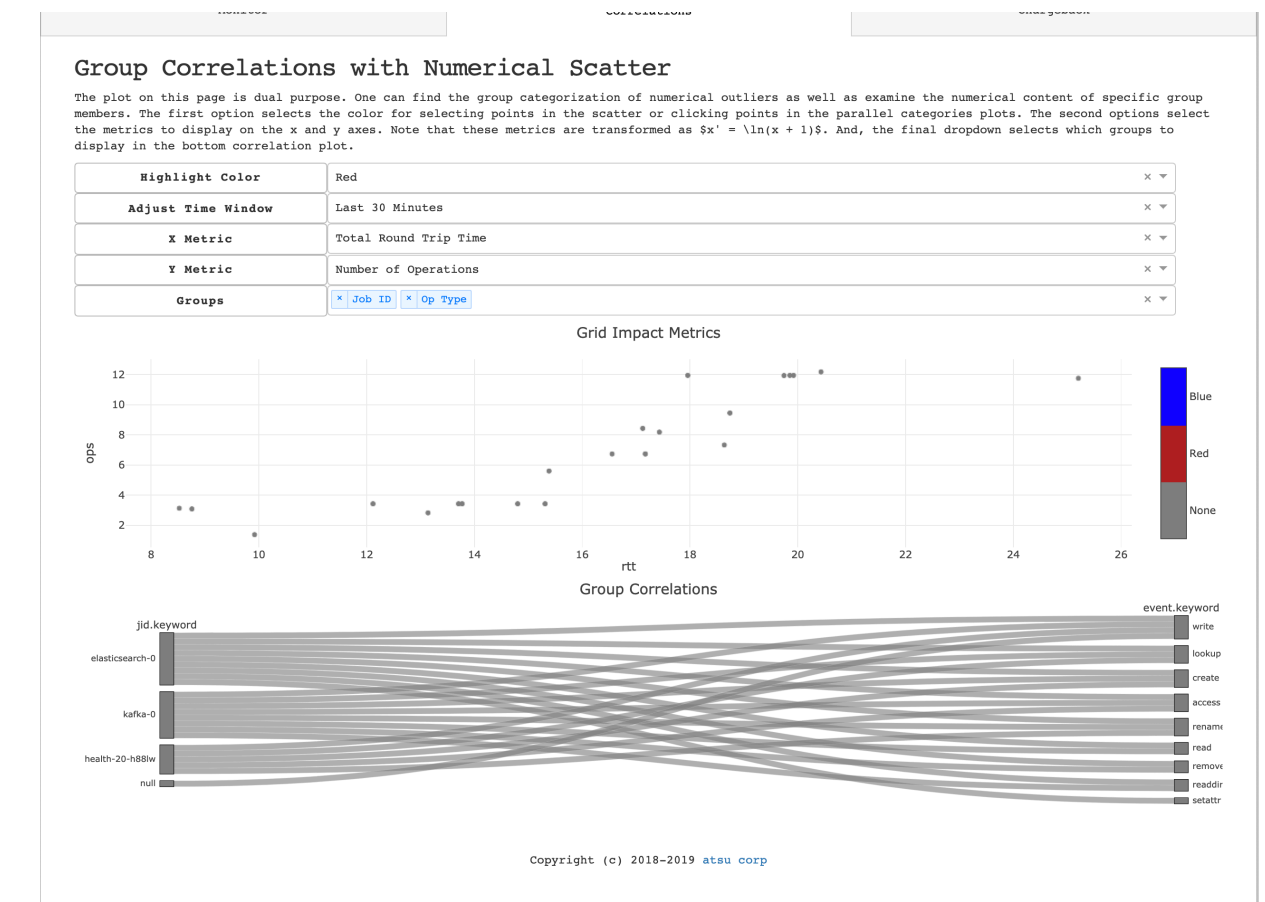
user/job IO chargeback



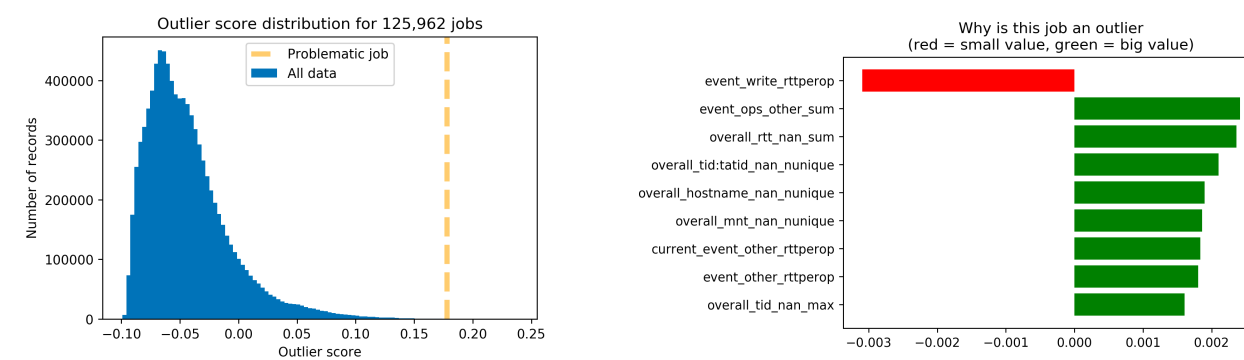
real-time NFS/SMB visibility



software/I/O usage




job anomaly detection

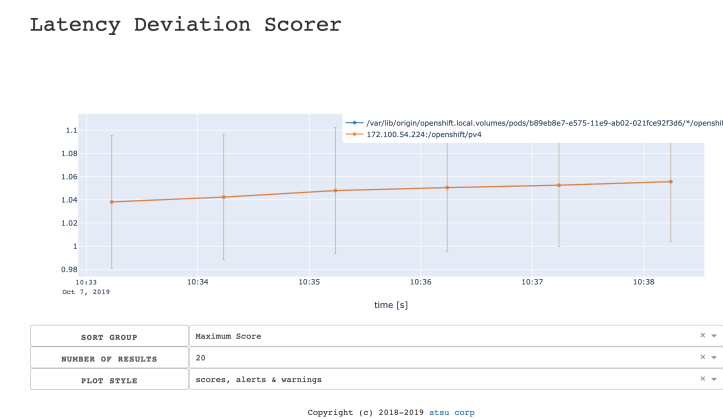


Job anomaly score explanation

- JobID: 128701194
- Anomaly Score: 1.1005
- Anomaly record ID: b4830e6c-6cdf-4305-adf9-a11e9b5bcb14



overwhelmed storage detection

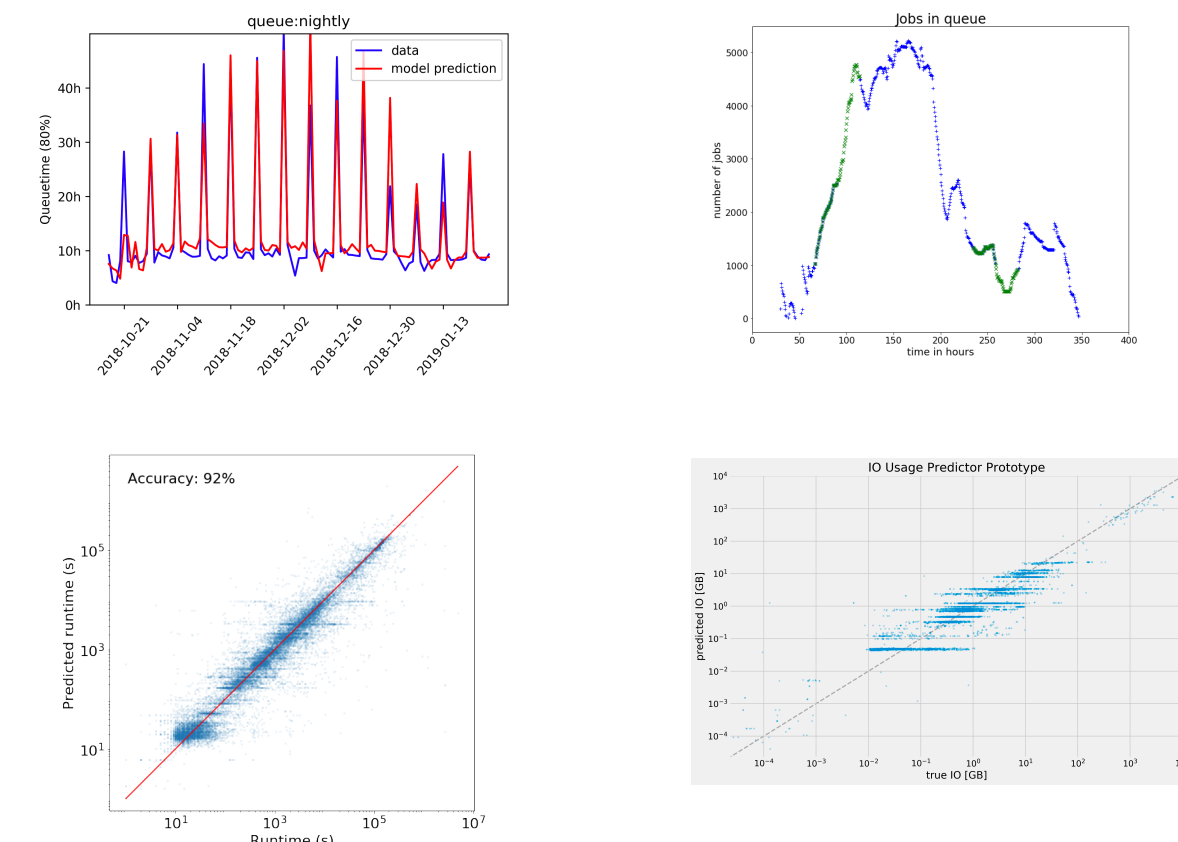


atsu-testing APP 10:32 AM
 rtt.deviation.alert alert ['field': 'mnt', 'value':
 '/var/lib/origin/openshift.local.volumes/pods/b89eb8e7-e575-11e9-ab02-021fce92f3d6/volumes/kubernetes.io~nfs/openshift-nfs-4'] with a value of
 5.630114165151673 > 1.2

View Alert - 15ee4f9d-1f04-487c-b95f-83c76fd0c7c2





queue and job prediction



An operation under examination

1




atsu agent sees NFS/SMB operation

- via kernel module or eBPF/ETW
- trivial overhead (μs)
- asynchronous to data bus

An operation under examination

1



atsu agent sees NFS/SMB operation

- via kernel module or eBPF/ETW
- trivial overhead (μs)
- asynchronous to data bus

2




atsu agent sends metadata to data bus

- minimal batch between operation and report
- multiple topics per types of data stream

An operation under examination

1



atsu agent sees NFS/SMB operation

- via kernel module or eBPF/ETW
- trivial overhead (μs)
- asynchronous to data bus

atsu agent sends metadata to data bus

2



- minimal batch between operation and report
- multiple topics per types of data stream

atsu service environment (hybrid SaaS)

3



- provides **detailed view**
- obfuscates data, transfers to atsu modeling
- streaming aggregations, predictions, etc.

atsu service

anomaly detection
resource prediction
submission suggestions



atsu.io

model creation
deployment

health



obfuscation


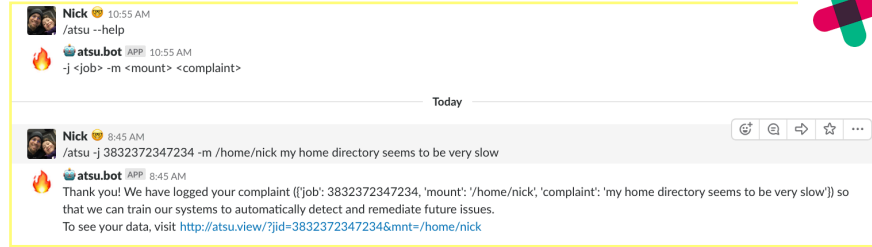
aggregator(s)
predictor(s)
job-control
queue-export
view

view



JID	total bytes	read bytes	write bytes	total ops	Mount	total bytes	read bytes	write bytes	total ops
khaidi-1-4qk1	456 MB	617 MB	38.8 MB	322424	/var/lib/origin/openshift.local.volumes/pods/...	1.23 GB	635 MB	421 MB	44993
elasticsearch-0	273 MB	2.35 MB	270 MB	114315					
katka-0	82.8 MB	15.0 MB	67.8 MB	5367					
health-44-hsh2	48.0 MB	0 B	48.0 MB	30.0					
snakemake-0	1.00 B	0 B	1.00 B	12.0					

File Path	Mount	total bytes	read bytes	write bytes	total ops
/logs/atsu-emp-gather...	/var/lib/origin/openshift.local.volumes/pods/...	44.2 MB	0 B	44.2 MB	1843
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	20.1 MB	0 B	20.1 MB	46.0
/logs/atsu-emp-gather...	/var/lib/origin/openshift.local.volumes/pods/...	16.4 MB	0 B	16.4 MB	495
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	9.48 MB	0 B	9.48 MB	472
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	9.24 MB	0 B	9.24 MB	23.0
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	9.11 MB	0 B	9.11 MB	491
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	8.79 MB	0 B	8.79 MB	850
/node/...	/var/lib/origin/openshift.local.volumes/pods/...	8.68 MB	0 B	8.68 MB	582
/logs/atsu-emp-gather...	/var/lib/origin/openshift.local.volumes/pods/...	7.55 MB	7.55 MB	252 MB	80.0
/logs/atsu-emp-gather...	/var/lib/origin/openshift.local.volumes/pods/...	7.52 MB	7.52 MB	108 MB	44.0




IO/CPU meta-data


Central processing block containing obfuscation, aggregator(s), predictor(s), job-control, queue-export, and view components.

atsu benefits


Anticipates potential problems
Identifies complex faults
Increases overall efficiency
Simplifies environment scaling



Predictable



Efficient




Agile

atsu benefits


Save money.

Go faster.


Get better automagically!



Predictable



Efficient



Agile


atsu service

On-premise


Hosted

Hybrid


Start your POC today!



Predictable




Efficient



Agile



atsu.io



Thanks!