

The Challenges of Managing Real-Time Financial Market Data Storage

SREcon23 Asia/Pacific
June 16, 2023

Kiran Kasichayanula, System Reliability Engineer
Nishith Nedungadi, SRE Team Leader
Feeds Engineering

[TechAtBloomberg.com](https://techatbloomberg.com)

© 2023 Bloomberg Finance L.P. All rights reserved.

Engineering

Bloomberg

Agenda

- Introduction to market data
- Existing workflows and challenges in managing data storage
- Deep dive into our solution
- The road ahead



The ***Bloomberg Terminal*** is software that delivers a diverse array of information, news and analytics to facilitate financial decision-making.


TechAtBloomberg.com

© 2023 Bloomberg Finance L.P. All rights reserved.

Bloomberg

Engineering

Primer: What is Market Data?



Instrument 4755 JP ¥ ↓ e 544.2 -7.8 544 / 544.3 159k x 2900

At 2:00 d Vol 31,384,000 0 548.5T H 558.6T L 544T Val 17.249B

755 JP Equity Settings

Exchanges: JF JS JT JN JV JE JI

Orders

E	Total	Size	Bid	Ask	Size	Total E
JT	158900	158900	544	544.3	2900	2900 JT
JT	171800	12900	543.9	544.4	3400	6300 JT
JT	193900	22100	543.8	544.5	1500	7800 JT
JT	204300	10400	543.7	544.6	1600	9400 JT
JT	217200	12900	543.6	544.7	1800	11200 JT

Trades

Time	Size	Price E
02:22:00	43.1k	544.2 JT
02:11:00	46.4k	544.2 JT
02:10:00	1.2k	544.2 JT
02:10:00	1k	544.2 JT
02:10:00	22k	544.2 JT
02:07:00	900	544.2 JT
23:04:00	4k	550.8 JT
22:56:00	2k	550.9 JT
22:44:00	100	550.9 JT
19:21:00	1k	552 JT
19:20:00	4k	552 JT
02:20:00	57.7k	544.2 JT

Price Book

Buy	Amount	Sell	Remaining
0	0	0	0

Avg Price 544.3000

Trade Recap | QR »

Volume Distribution | QM »

Rakuten Group Inc Trade Lot Size 100
Change on Day -7.8 Avg Vol 5 Day 53,090,700
% Change -1.41% VWAP 549.6301
Open 548.5 Maximum Limit Price
High 558.6 Low 544 Minimum Limit Price

Primer: What is Market Data?

Normalization: The representation of information as basic as trade price and trade quantity can vary substantially across major exchanges.

Price and quantity have different data types

	Tag	Tag Name	Req	Data Type	Comment	EQT	DER
→	270	MDEntryPx	C	Price	Price of the Market Data Entry. Required when this market data entry involves a price. Represents the notional value for trade volume (B). Other entry types that do not involve price do not require this tag.	X	X
→	271	MDEntrySize	C	Qty	Quantity or volume represented by the Market Data Entry. Required when MDUpdateAction = New (0) and MDEntryType = Bid (0), Offer (1), Trade (2), Trade Volume (B) or Opening Price (4).	X	X

Ultra fast, tightly packed price and volume with venue identification

Trade Price	2	Short
Trade Volume	2	Short
Primary Listing Market Participant ID	1	Char

Feed Handler Application

Normalization: The representation of information as basic as trade price and trade quantity can vary substantially across major exchanges.

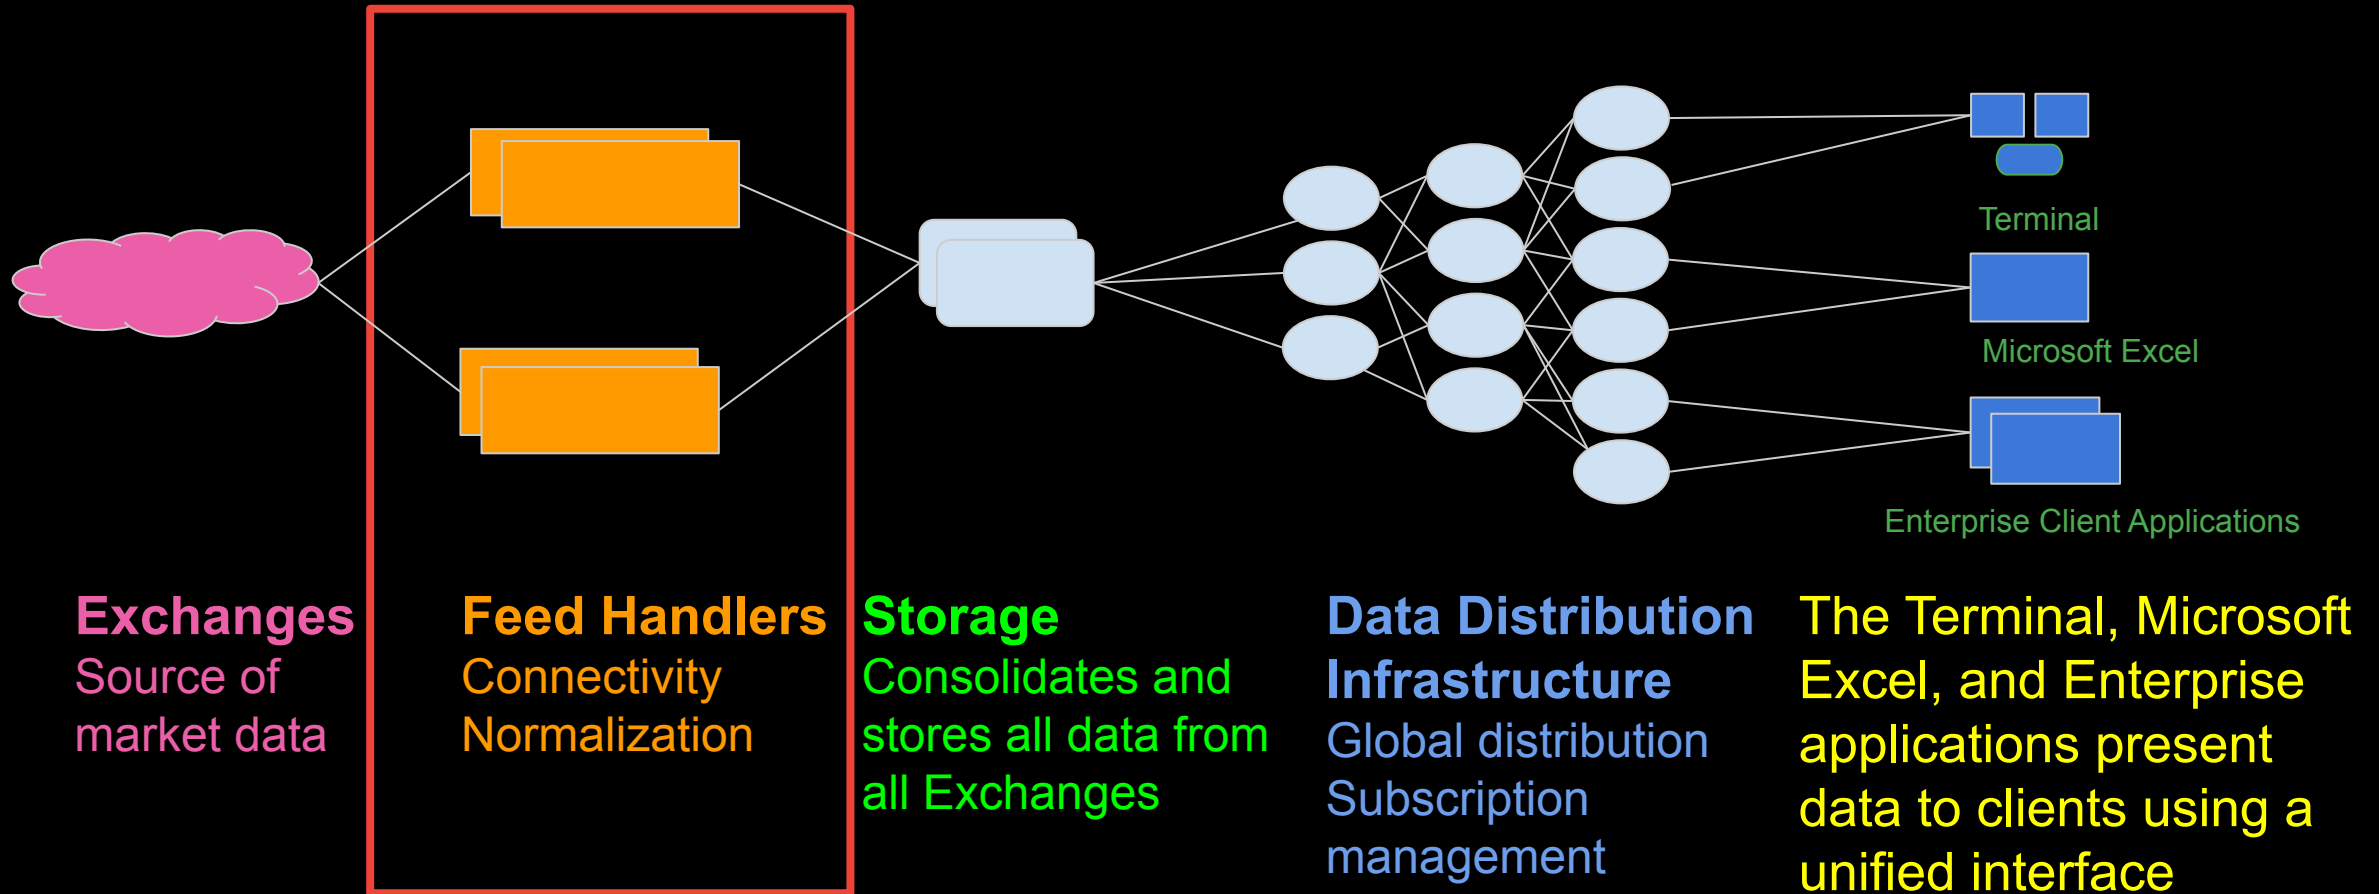
Tag	Tag Name	Req	Data Type
→ 270	MDEntryPx	C	Price
→ 271	MDEntrySize	C	Qty

Trade Price	2	Short
Trade Volume	2	Short
Primary Listing Market Participant ID	1	Char

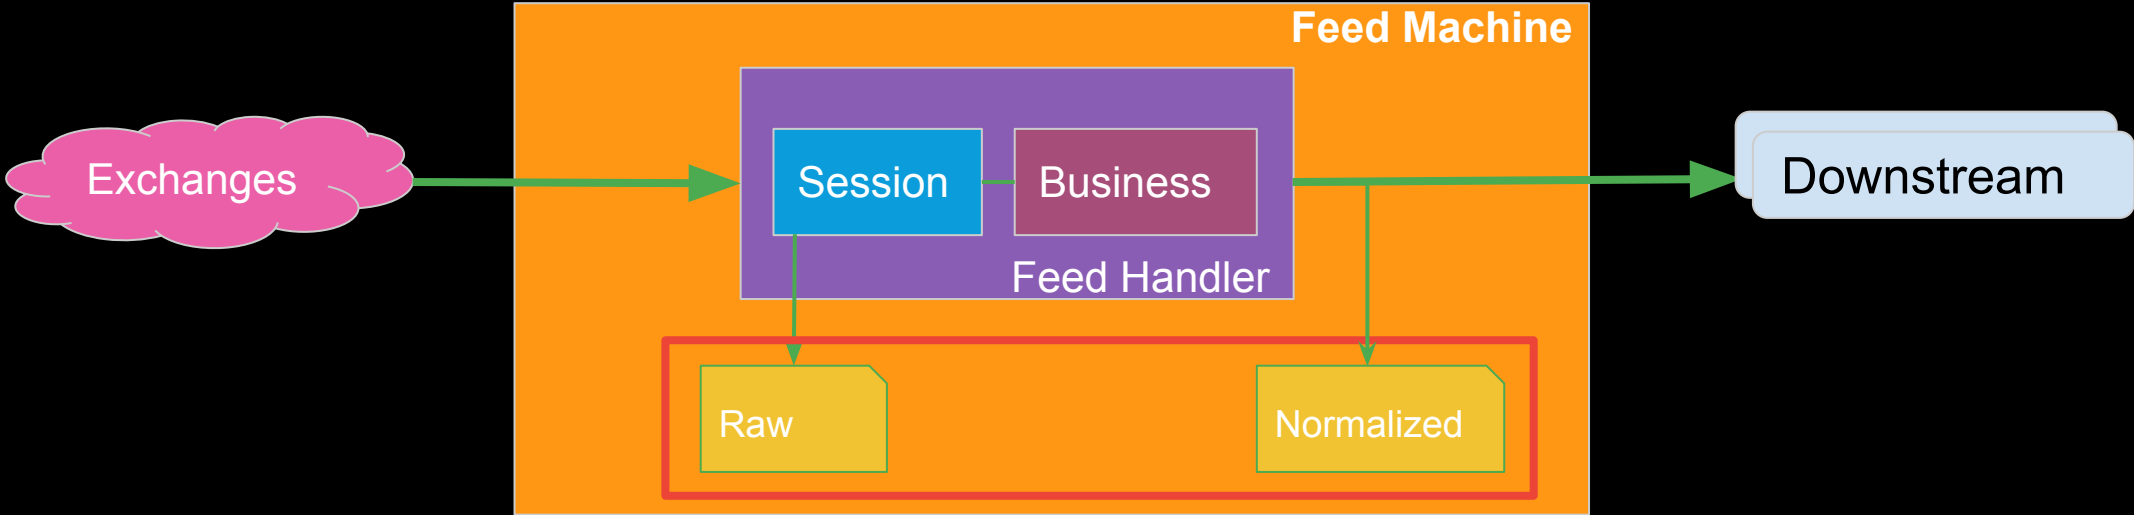
Example **Normalized** Market Data:

```
TradeWithTradeID 102|PriceClass  
1|FeSubSequence1 0|FeSubSequence2  
0|Sequence 952101|Timestamp 06/08/2023  
06:59:28|Key TEST |TradePrice 10 |TradeSize  
100|  
ConditionCodes 1: 19|TradeID 0001|
```


Bloomberg Market Data Pipeline

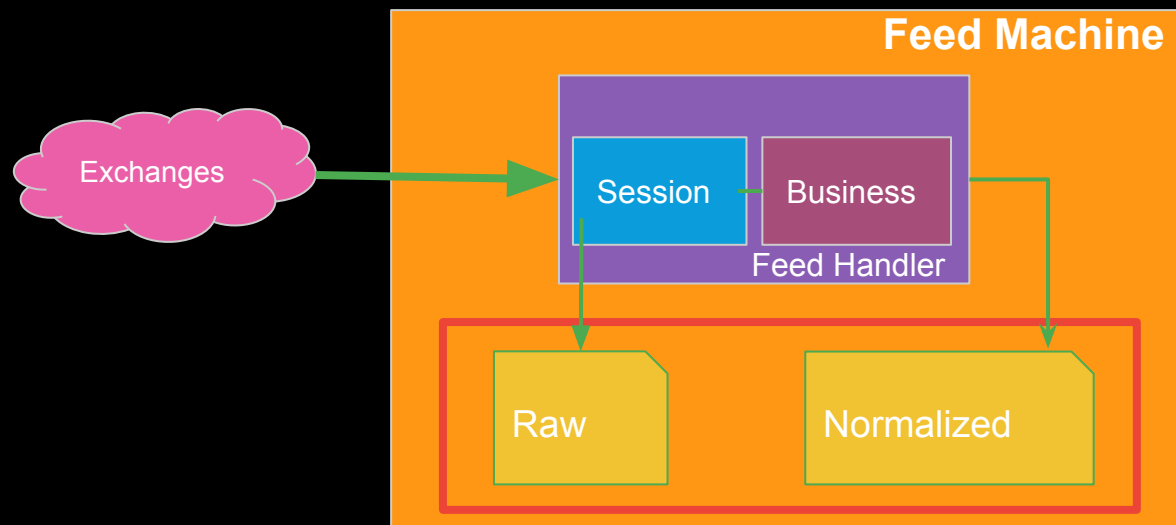


Feed Handler Application



Challenge: Managing Raw / Normalized Data Storage

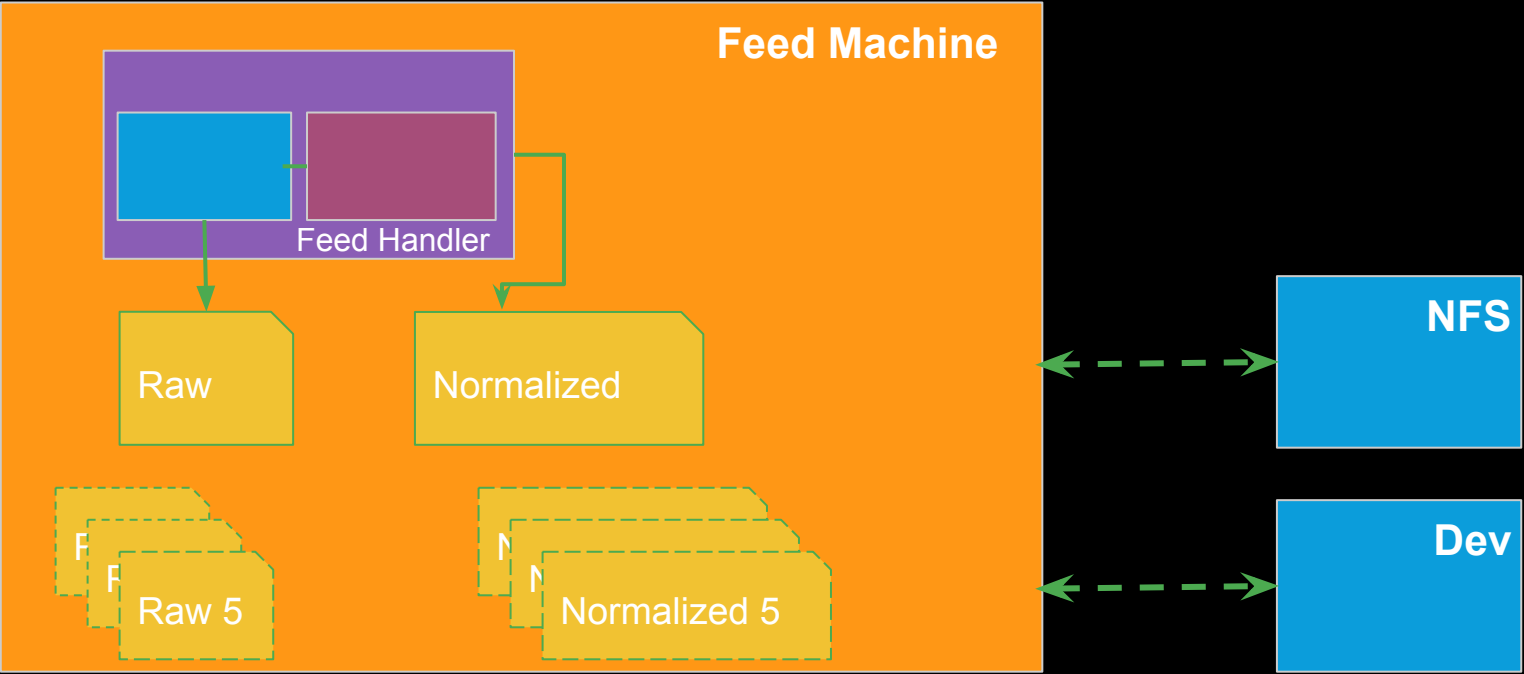
Business analysts and Engineering need a system to access the incoming raw market data and the normalized market data generated by the feed handlers



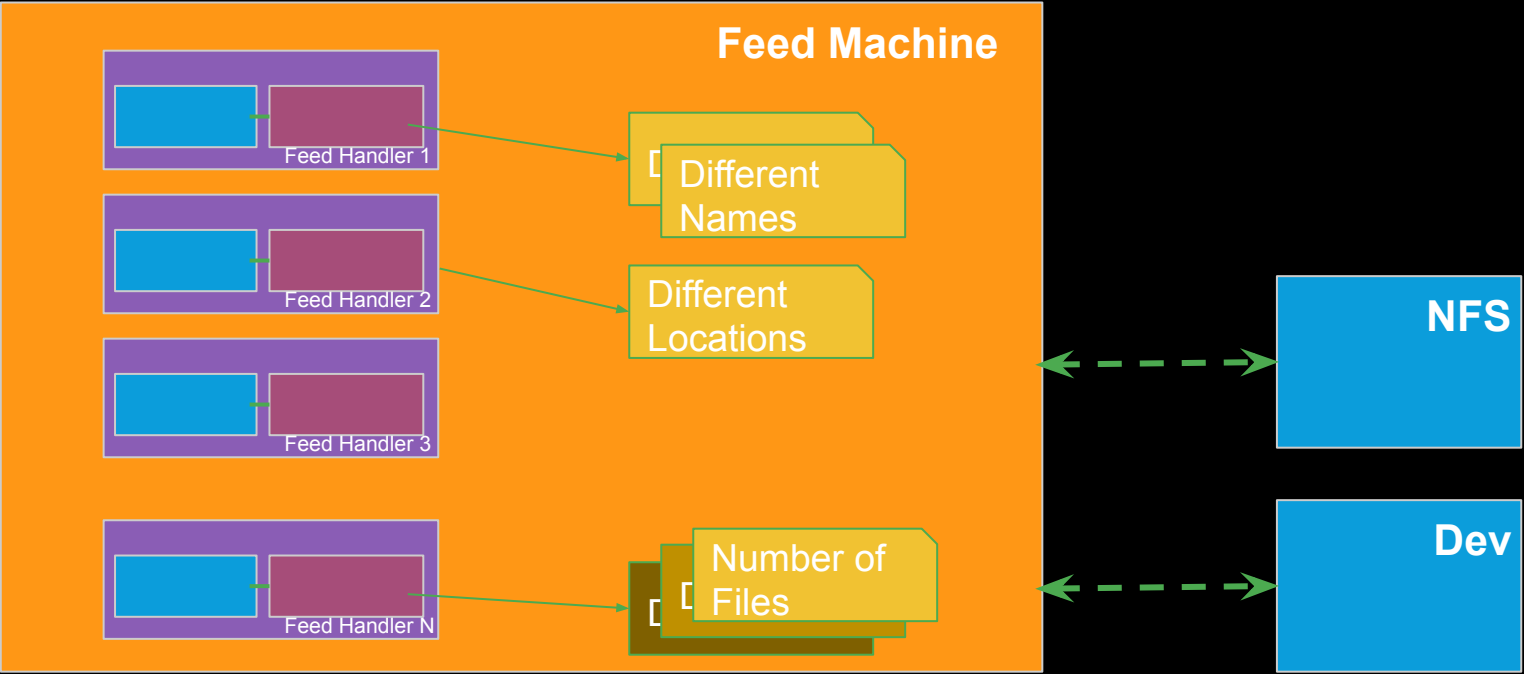
WHY?

- Business uses this data to understand and analyze raw data sent by exchanges
- To answer client queries
- Communication tool between Engineering and Business
- Debugging their applications
- To create new tests of interesting scenarios
- To validate the correctness of their processing
- Recover from outages

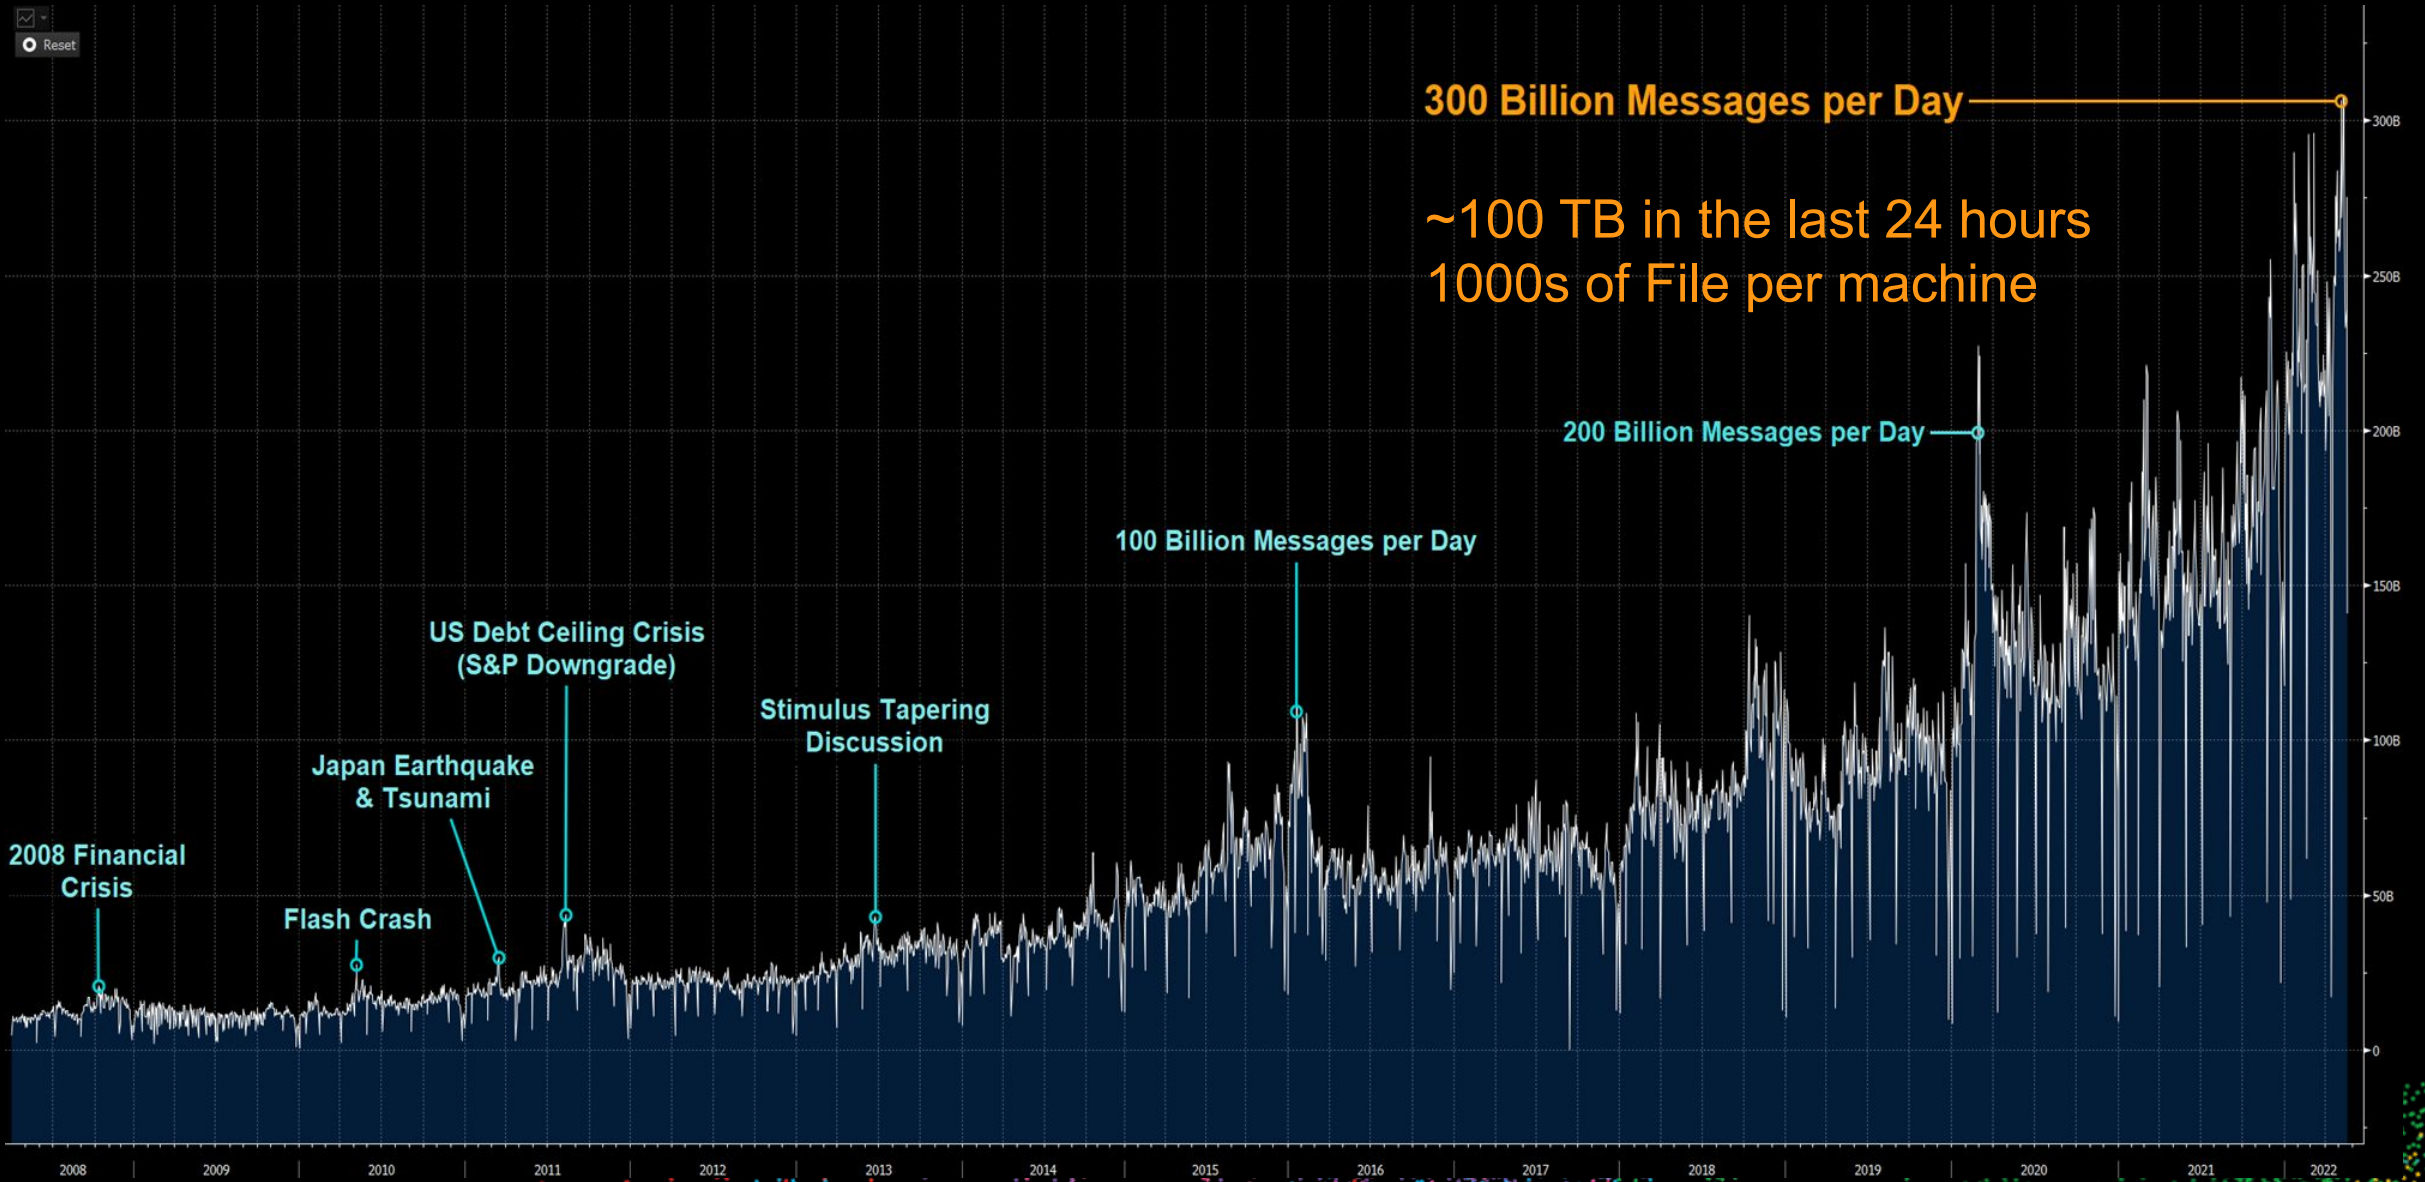
Managing Storage: Existing System



Managing Storage: Non-uniformity



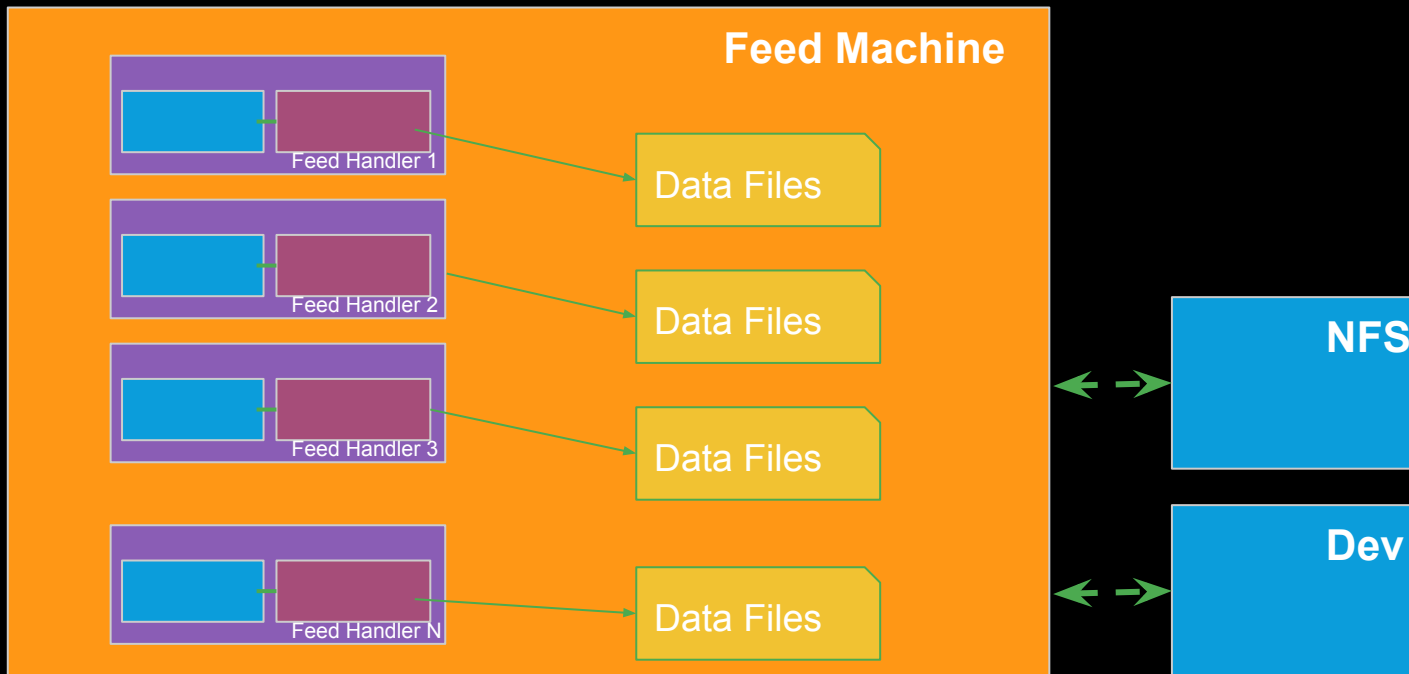
Managing Storage: Capacity impacts Uniform Policy



Managing Storage: Network zones security / bandwidth



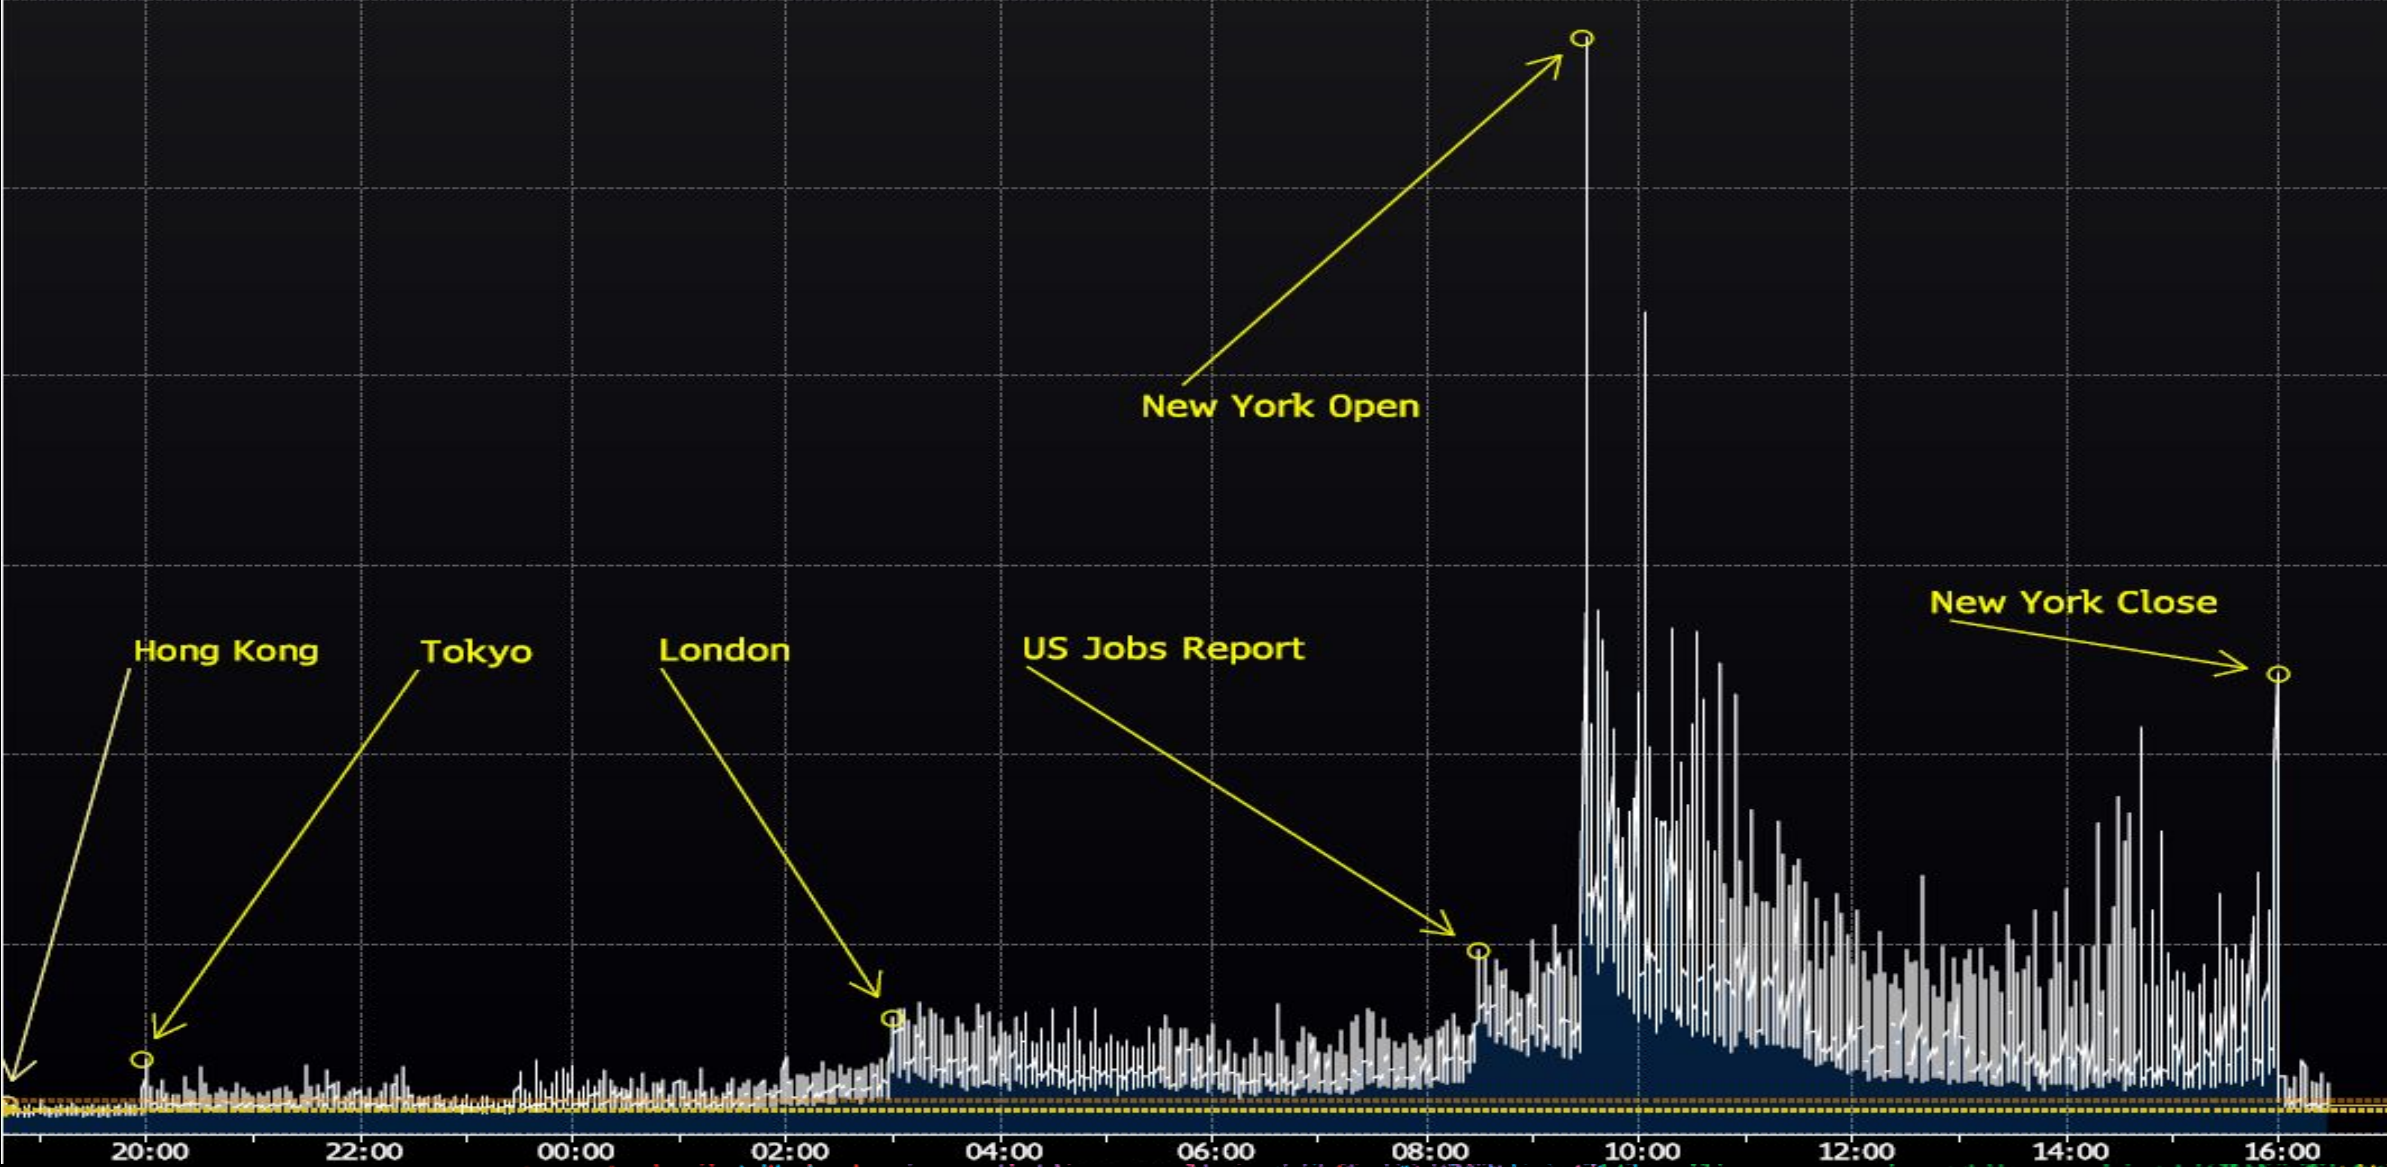
Managing Storage: Why Files?



WHY FILES?

- Existing interface
- Characteristic of market data

Market Data Rates: Spikiness



Requirement

Business analysts and Engineering need a system to access the incoming raw data and the normalized market data generated by the feed handlers

What we want

- Scalable
- Agile
- Secure
- Storage agnostic
- Datacenter agnostic
- Observable
- Correct
- Accessible



Architecture Deep Dive

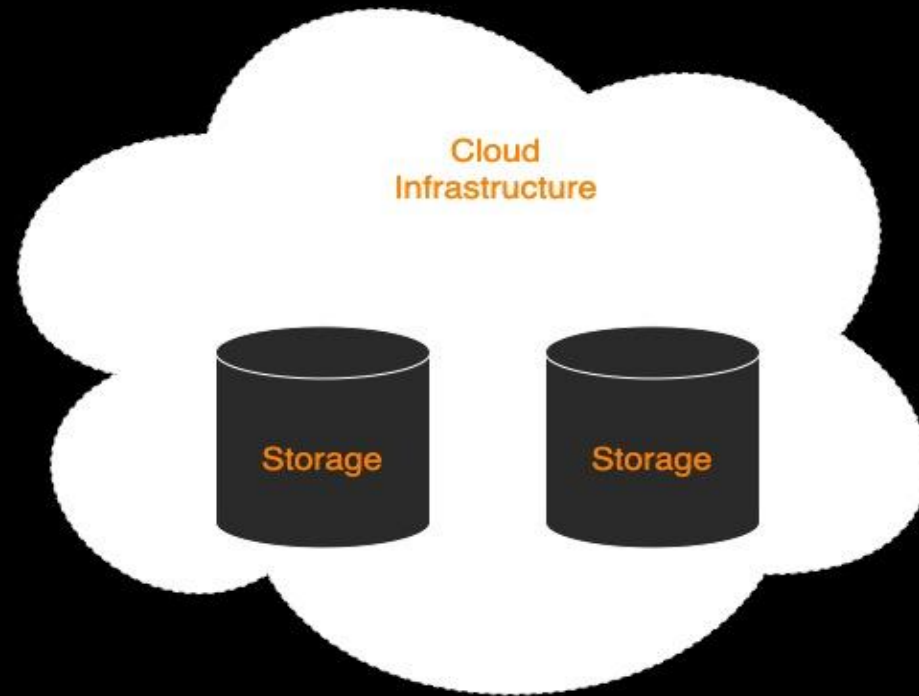
TechAtBloomberg.com

© 2023 Bloomberg Finance L.P. All rights reserved.

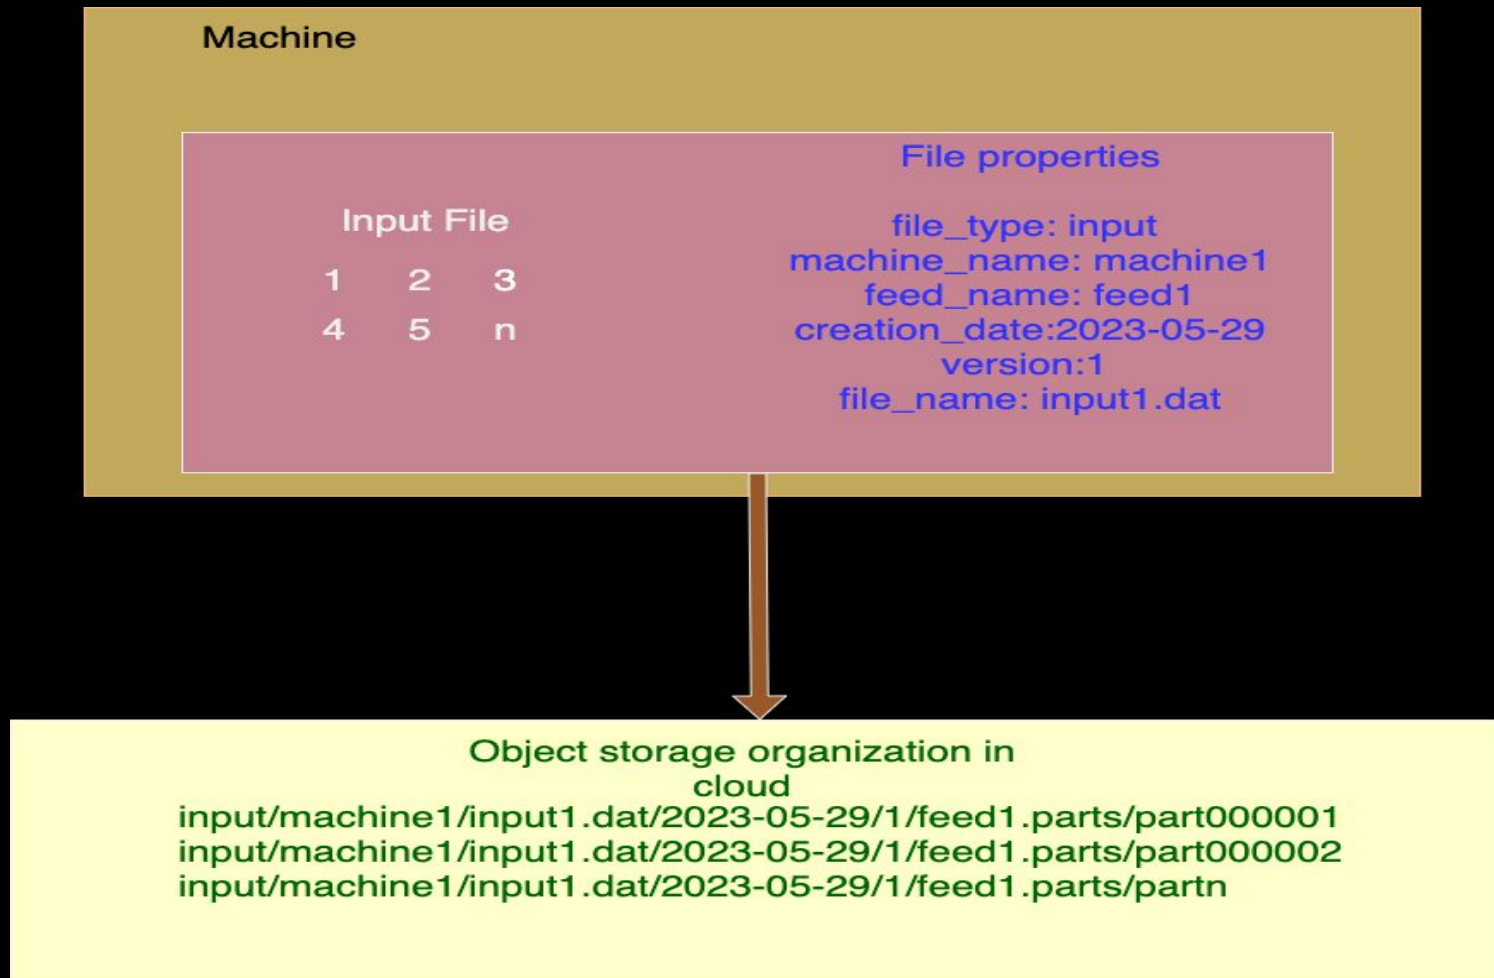
Bloomberg

Engineering


Architecture: Object Data Store in the Cloud



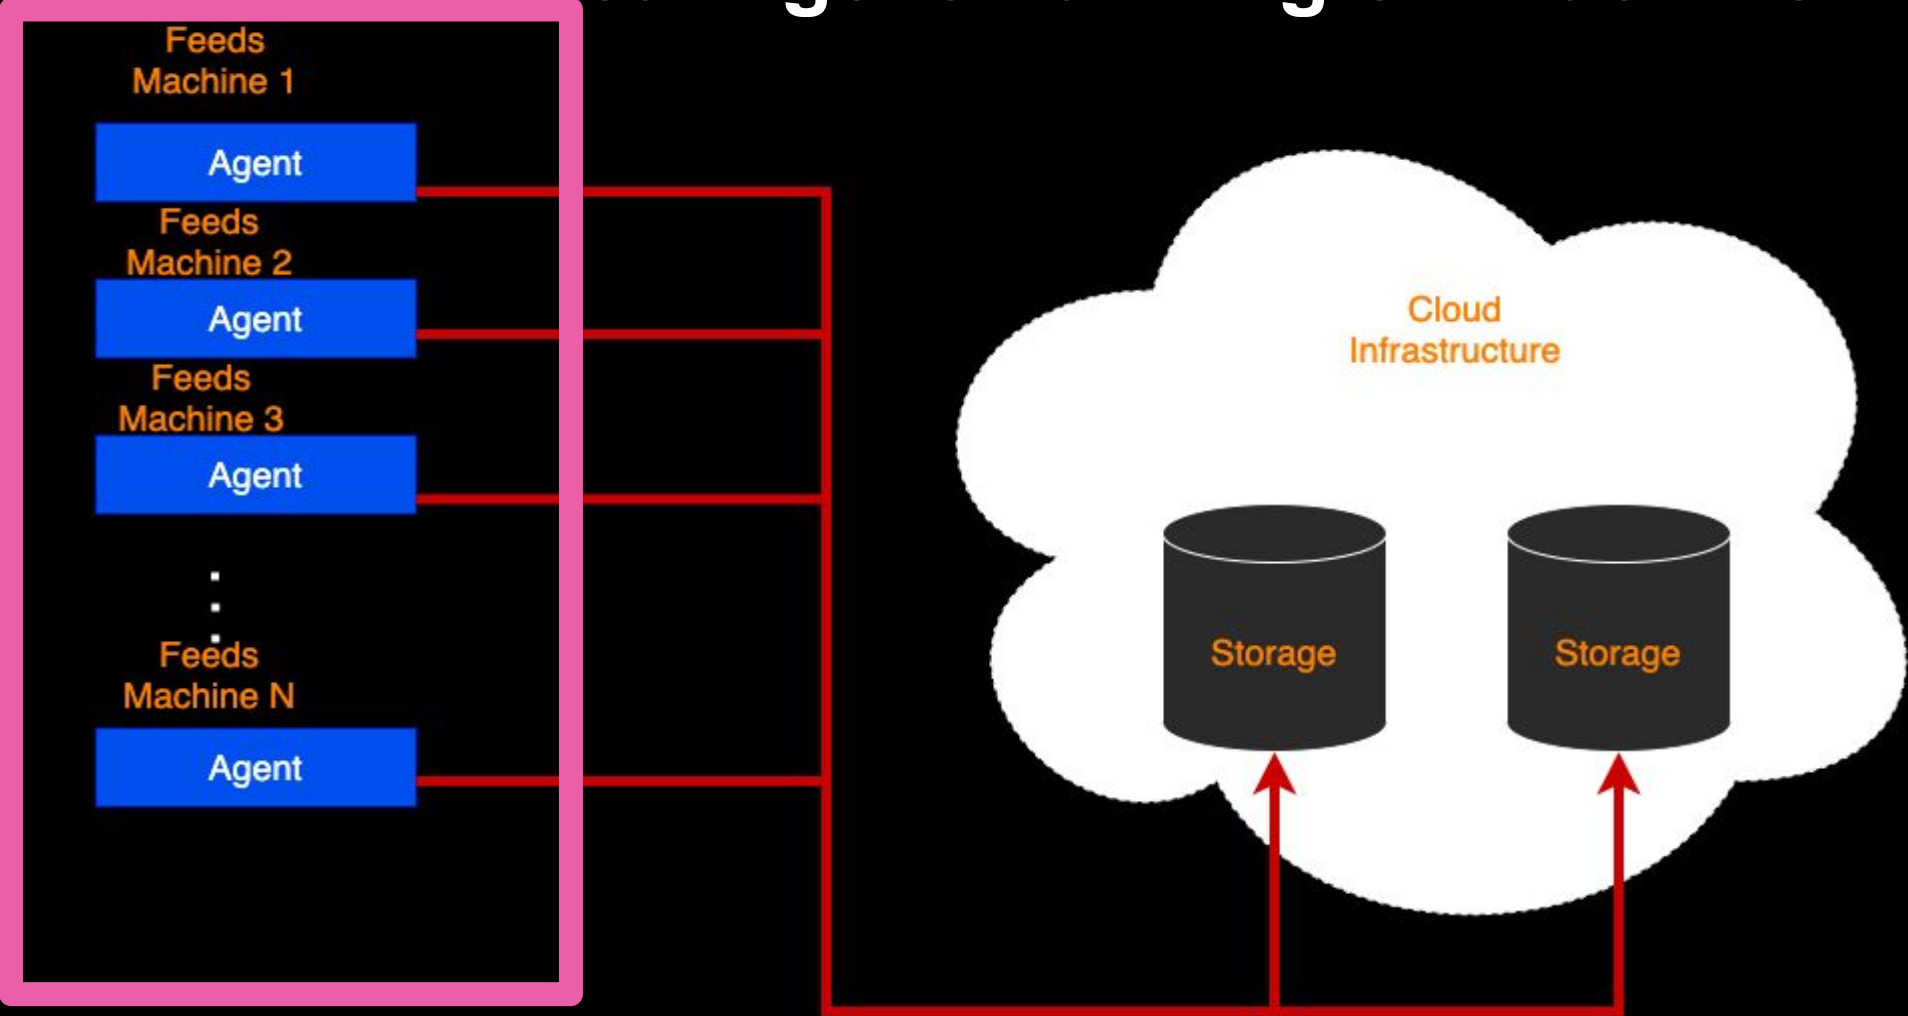
Architecture: Object Data Store in the Cloud



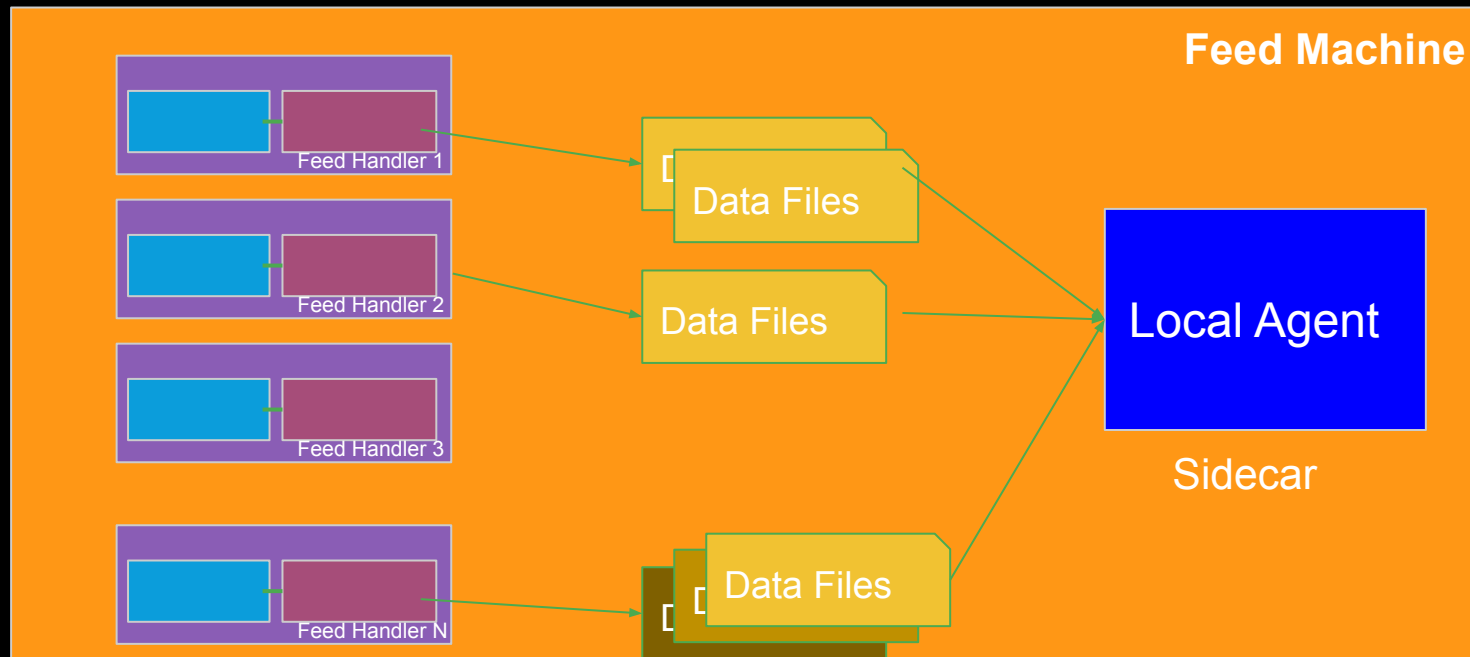
Architecture: Network Zones in the Cloud



Architecture: Local Agent Running on Machine



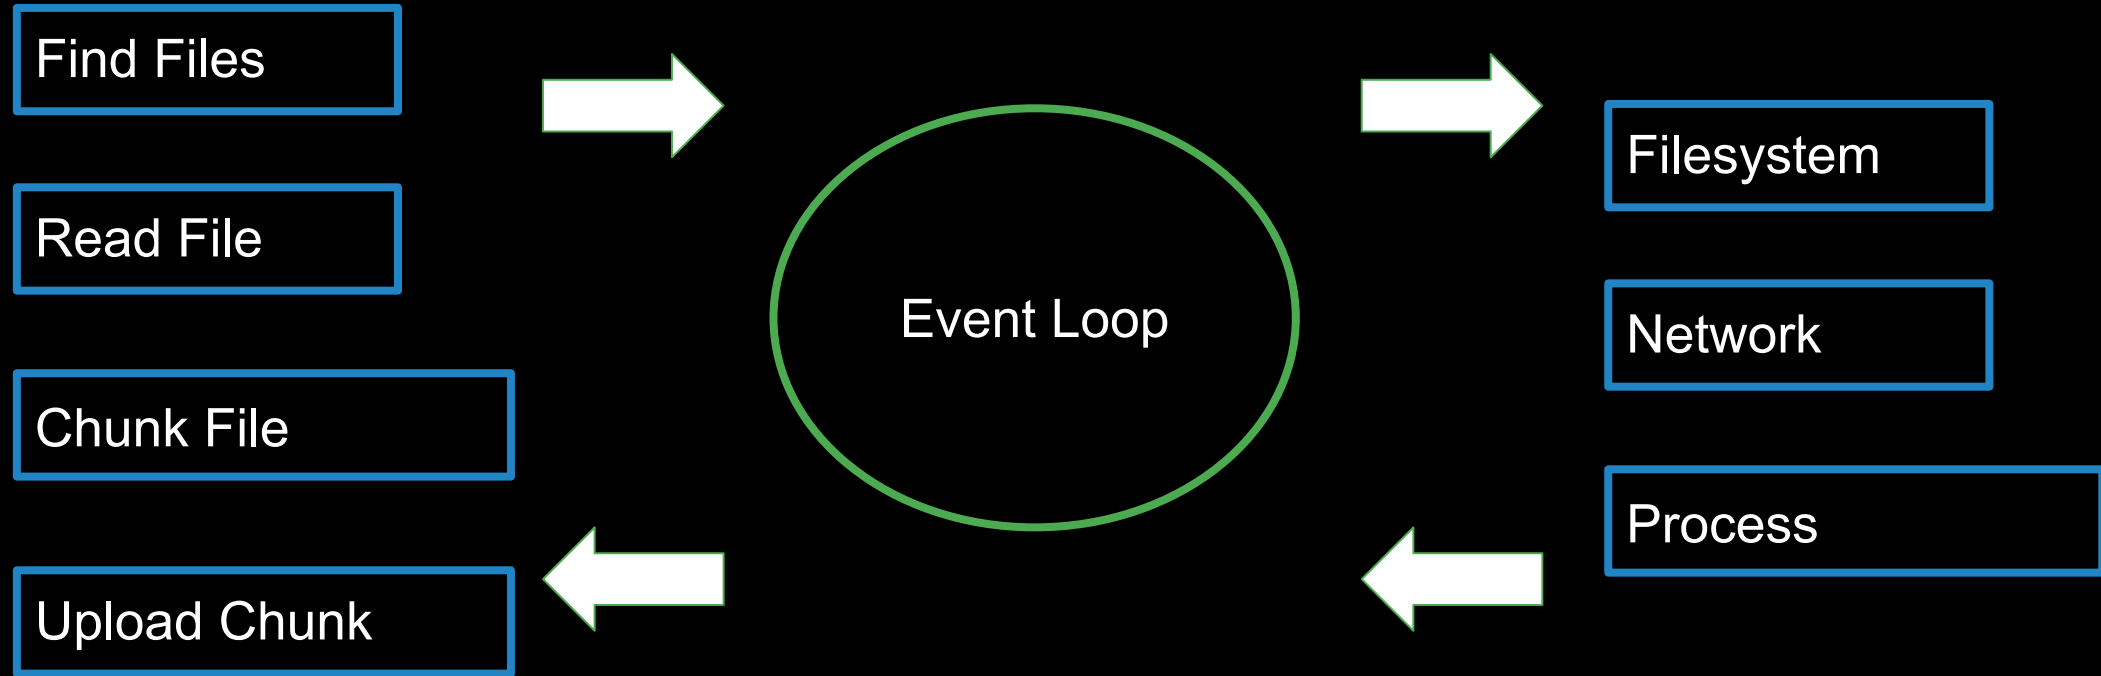
Architecture: Local Agent Running on Machine




Sidecar Properties :

Resource constrained
Appropriate NICE values


Architecture: Asyncio Library



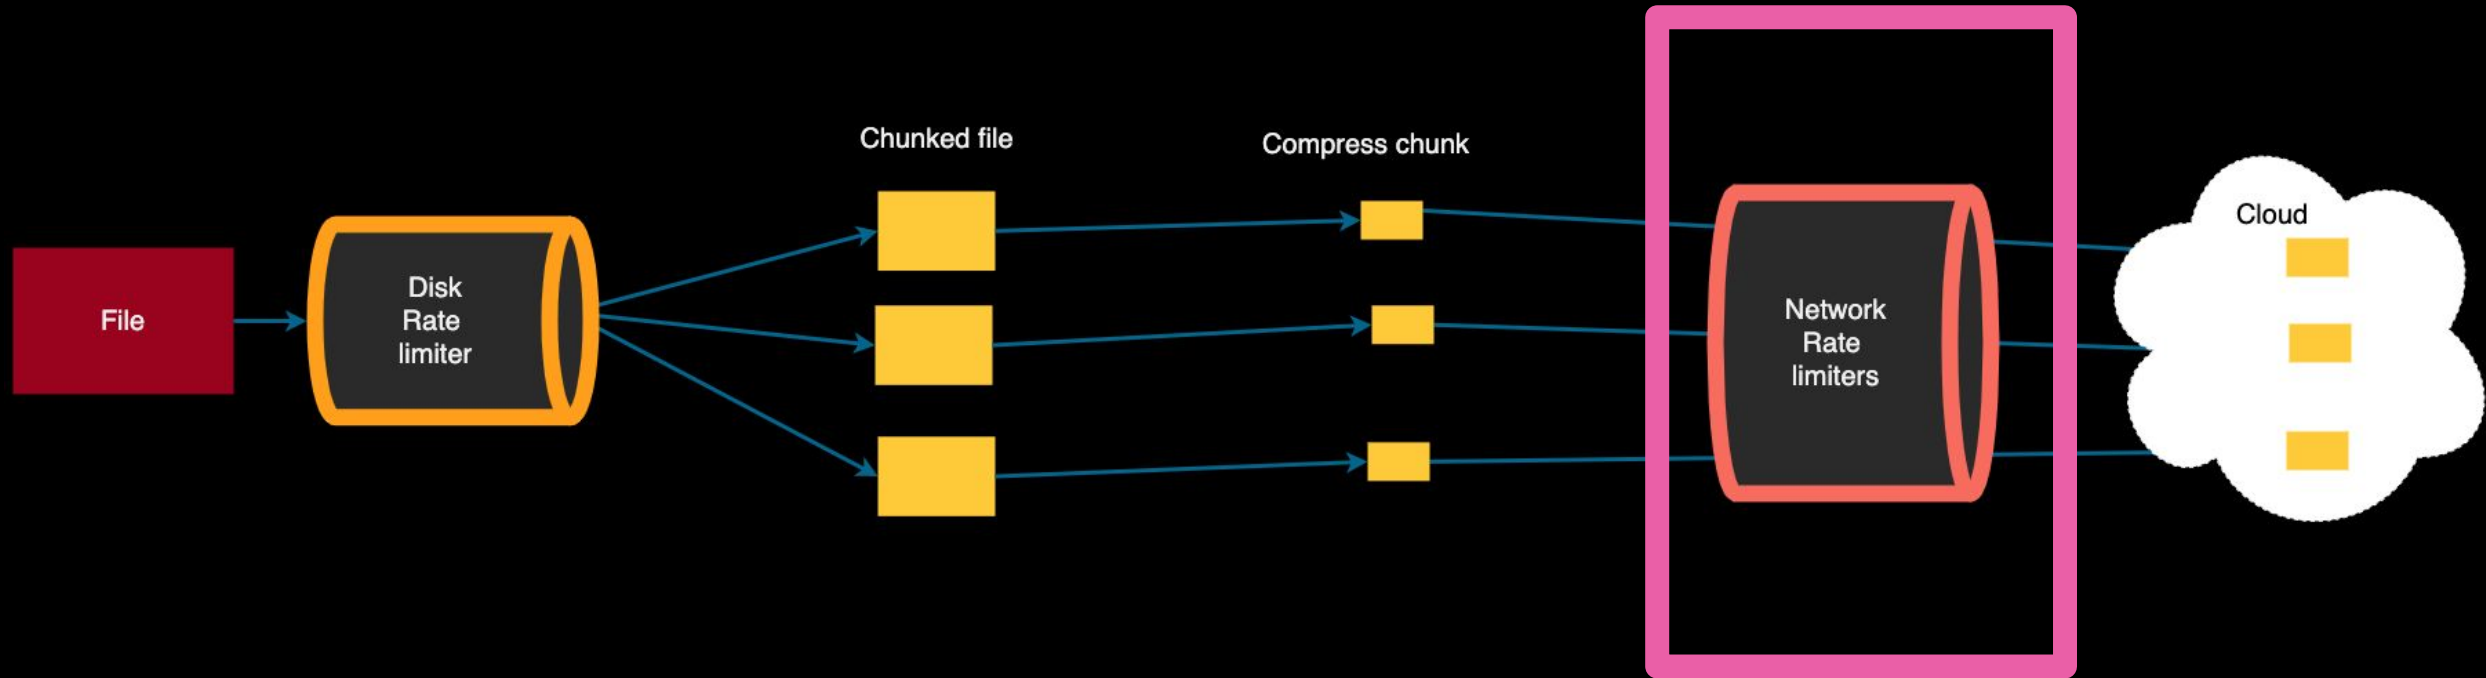
Architecture: Local Agent Running on Machine



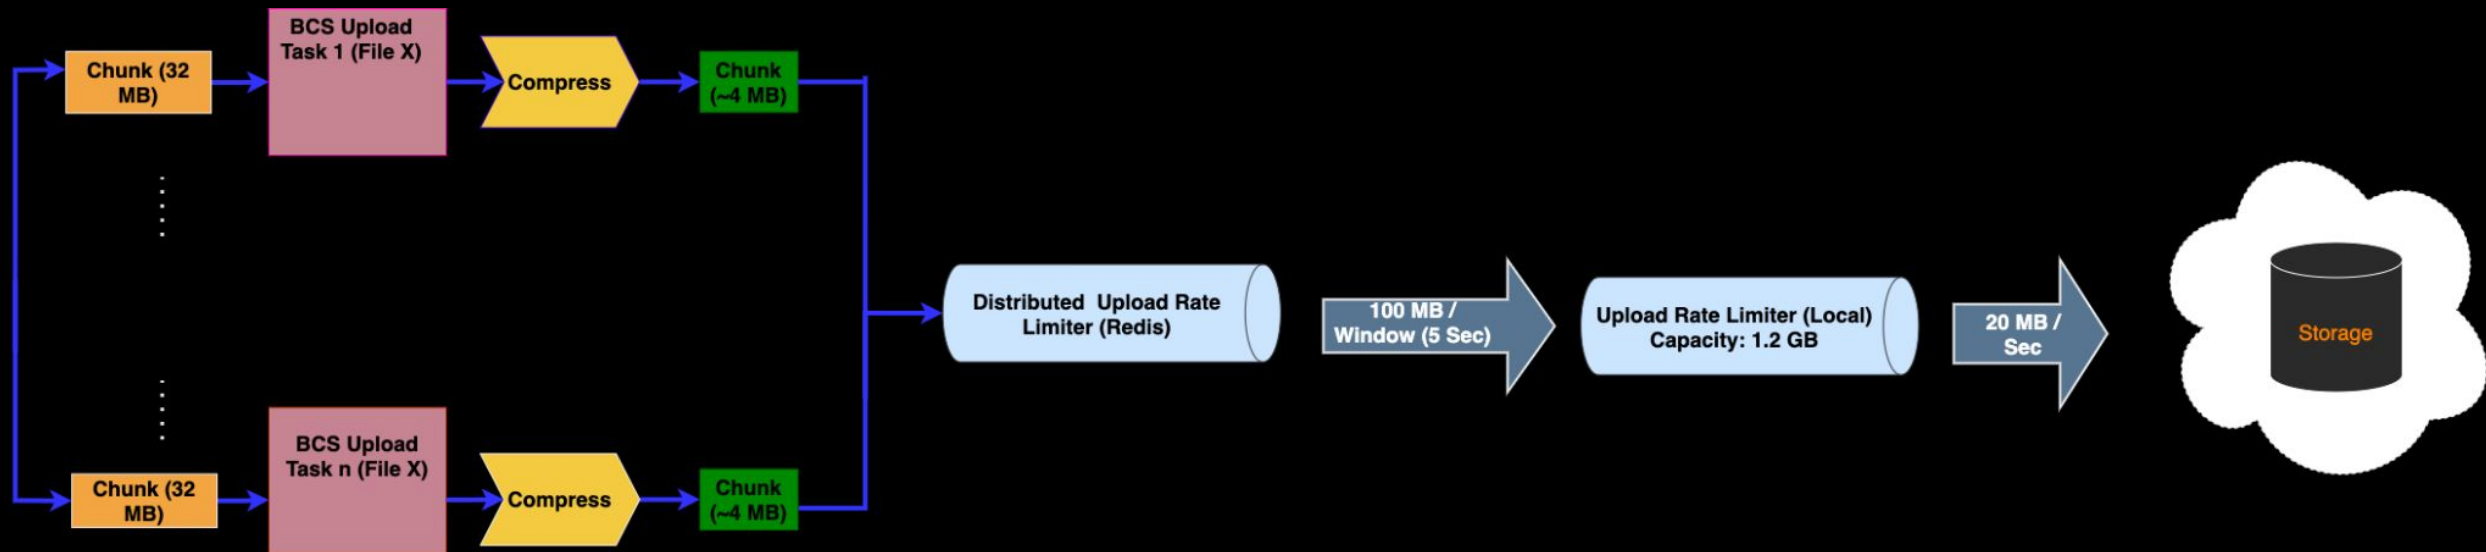
Architecture: Disk Rate Limiter



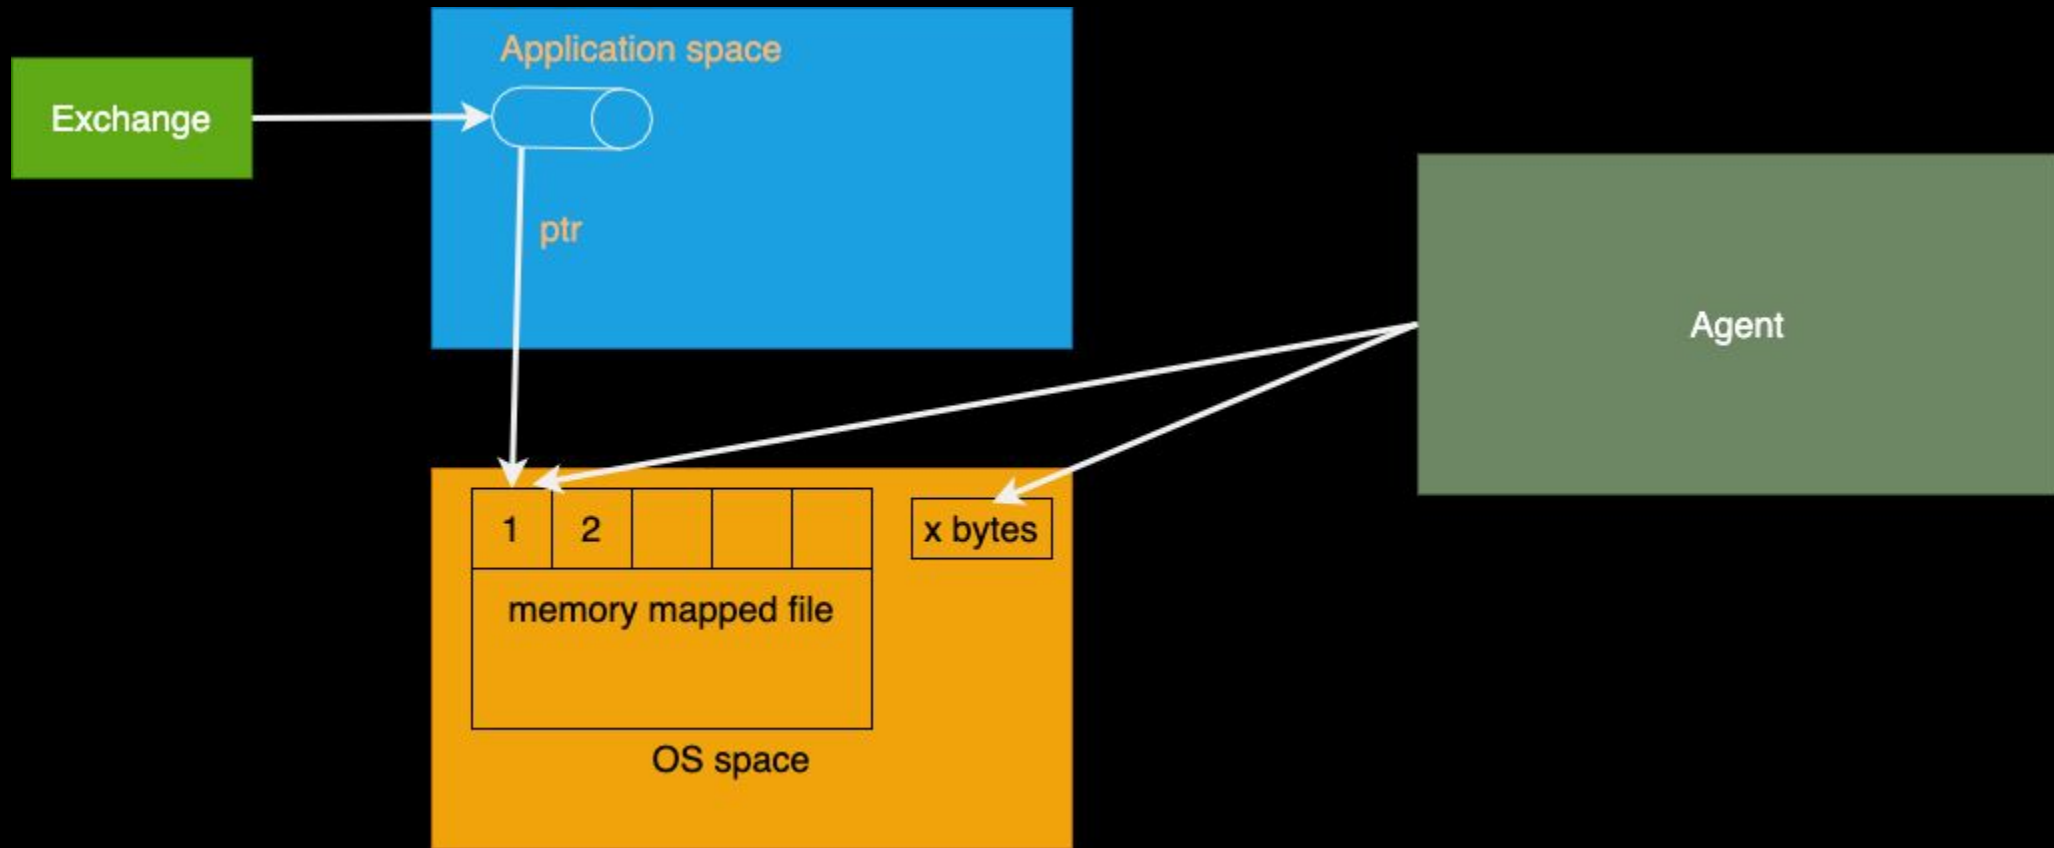
Architecture: Local Agent Running on Machine



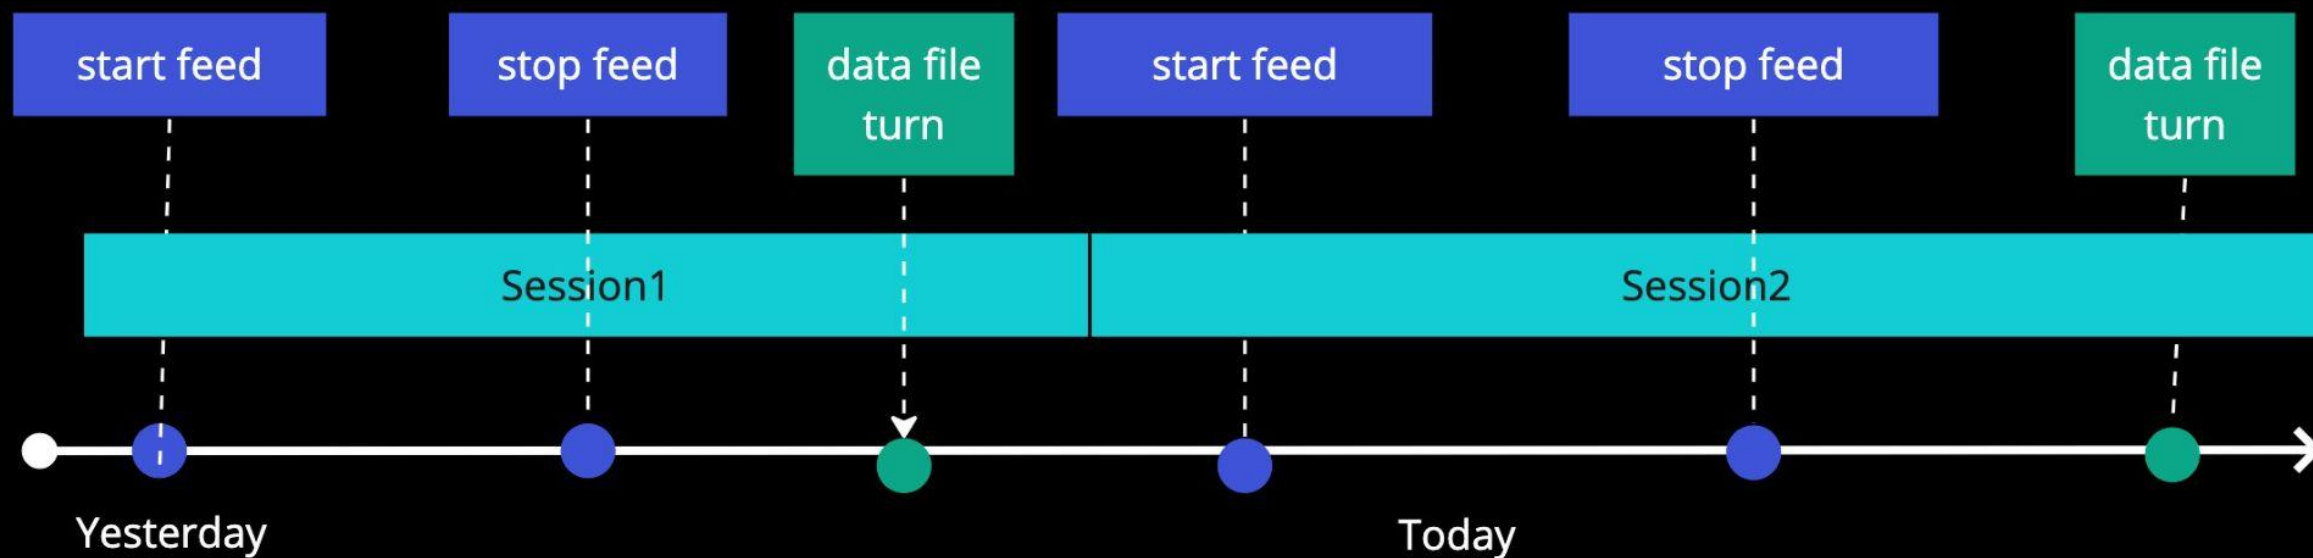
Network Rate Limiter




Application Aware



Application aware



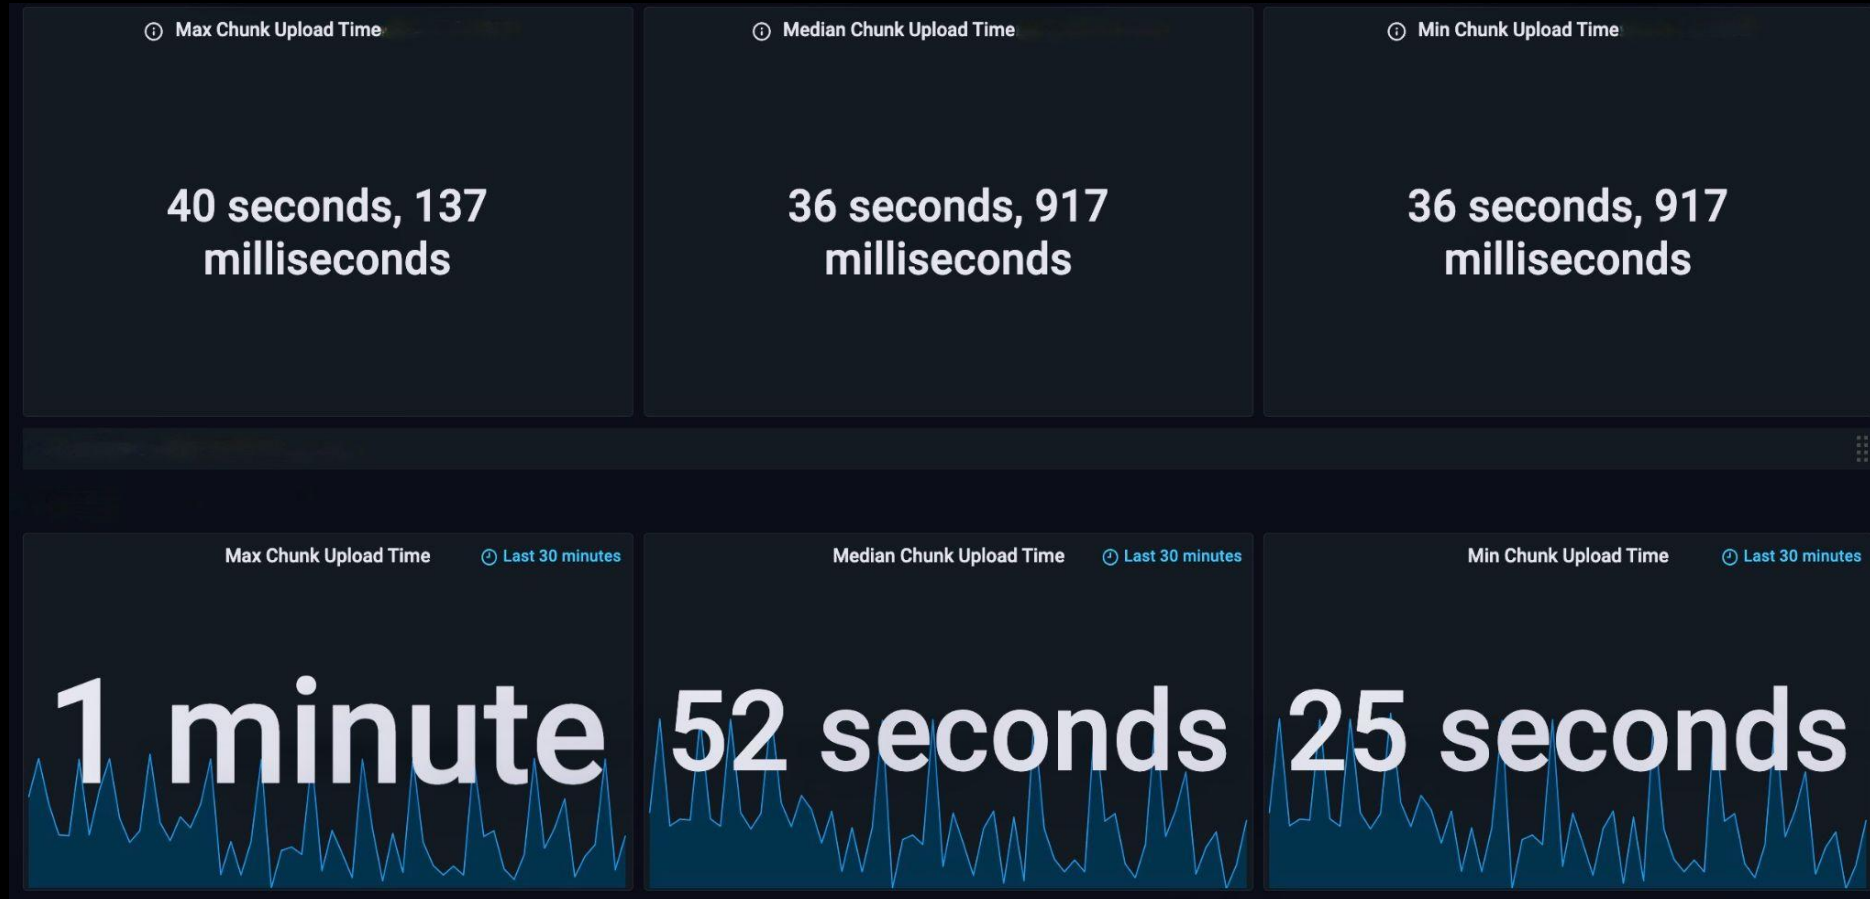
Architecture: Client APIs / Utilities



Client Access (Internal Business / Engineering)

- Business users (Web API powered by Kubernetes clusters)
 - On-demand import requests
 - Intra-session data
 - Centralized decoding
- Developers (Development machines)
 - Match input and output
 - Debug intra-day issues
 - Testing new parser deployment


Monitoring and dashboards



Bloomberg

Engineering

Architecture



The road ahead

- Scalability
- Real-time
- Improving observability

Takeaways

- A glimpse into real-world market data management challenges
- Applying cloud storage
- Considerations in building out a system

Bloomberg

Engineering

Questions?

[TechAtBloomberg.com](https://www.techatbloomberg.com)

© 2023 Bloomberg Finance L.P. All rights reserved.

Visit our booth

<https://TechAtBloomberg.com>

<https://www.bloomberg.com/careers>

Contact us:

kkasichayanu@bloomberg.net | nnedungadi@bloomberg.net

Engineering

Bloomberg

TechAtBloomberg.com