

SLA-constrained Feedback-based Software Load Distribution Algorithm that Minimizes Computing Resource Requirement

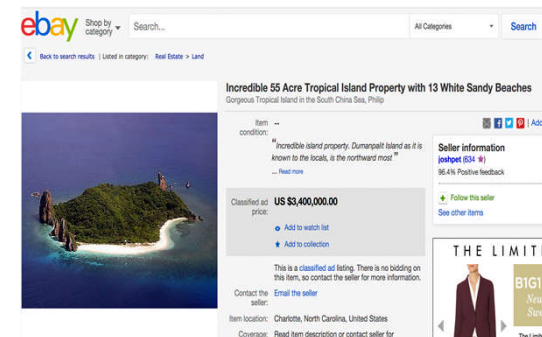
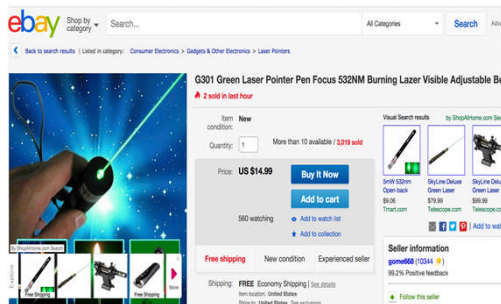
S.R. Venkatramanan (PayPal)
Rema Hariharan and Ashok Murthy (eBay)

Agenda

- Business Background
- Conventional Routing in the Cloud
- New algorithm
 - Principle
 - Description
 - Implementation Results
- Discussion



- The world's online marketplace
- eBay is a global commerce leader; connects millions of buyers and sellers around the world
- Facilitators for online sales.
 - Sold items range from broken pens to multi-million dollar homes and more.
 - Auction + Fixed price
- One of the largest data warehouses.



eBay BUSINESS



eBay Inc at a glance



\$2.1B

Revenue in Q1 2016



162M

Global Active Buyers



\$9B

Mobile Volume



57%

International
revenue



314M

App downloads



\$20.5B

GMV in Q1 2016



Money is changing



Mass Adoption of
Mobile Devices



Digitization
of Cash



Transformation
of Cards



Fragmentation of
Payment Types,
Technology and
Channels



Rise of Fraud
and Cybercrime

PayPal is leading the transformation



AT SCALE*



WITH MOMENTUM*



184 Million Active Accounts**



+17 Million Active Accounts Gained in 2015



\$282 Billion Total Payment Volume (\$66Billion Mobile)



+27% y/y TPV Growth****



\$9.24 Billion Revenue***



+24% y/y Transaction Growth



4.9 Billion Transactions (1.4B Mobile)

*Stats are full year 2015, unless otherwise noted.

**Stat is Q1 2015

***Revenue is presented on a non-GAAP pro forma basis, and includes the impact of pro forma adjustments directly attributable to the separation of the company from eBay Inc. on July 17, 2015 had they existed historically.

****Calculated on an FX-neutral basis

Our international footprint

79%

Of shoppers in key markets, 79% used PayPal for a cross-border transaction between June 2012 and June 2013.

94M

In 2014, 94 million people shopped with PayPal across borders in key markets.

\$300B

In key markets, cross-border trade is expected to pass \$300B in 2018.

Source: Modern Spice Routes, PayPal and Nielsen, July 2013
https://stories.paypal0corp.com/uploads/4/8/9/8/48984695/paypal_modernspiceroutes_report_final.pdf
Key Markets: US, UK, Germany, Brazil, China and Australia



Key Statistics

Developer Statistics



3000
engineers



x
10M
Lines



x 1000
Releases
/year



x 1000
Builds/d
ay



x 1000
Deploy/day


OpenStack Cloud



3 Regions



9 Availability
Zones



x 100K
Cores



> 10K
Physical
Servers



x 10K
VMs



x 10 PB
Storage

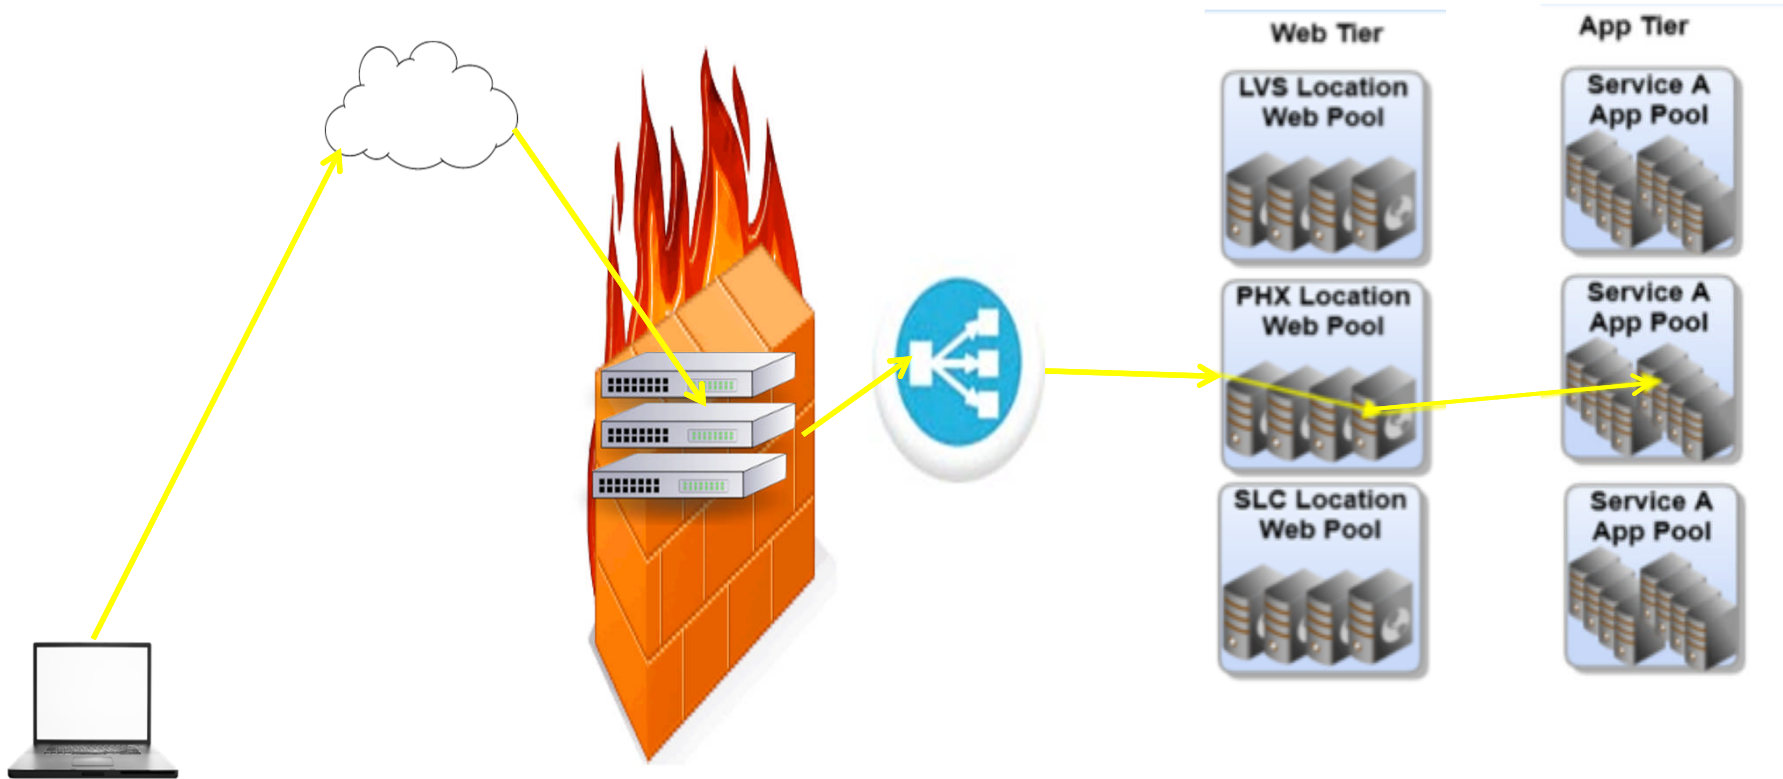


> 1000
Services

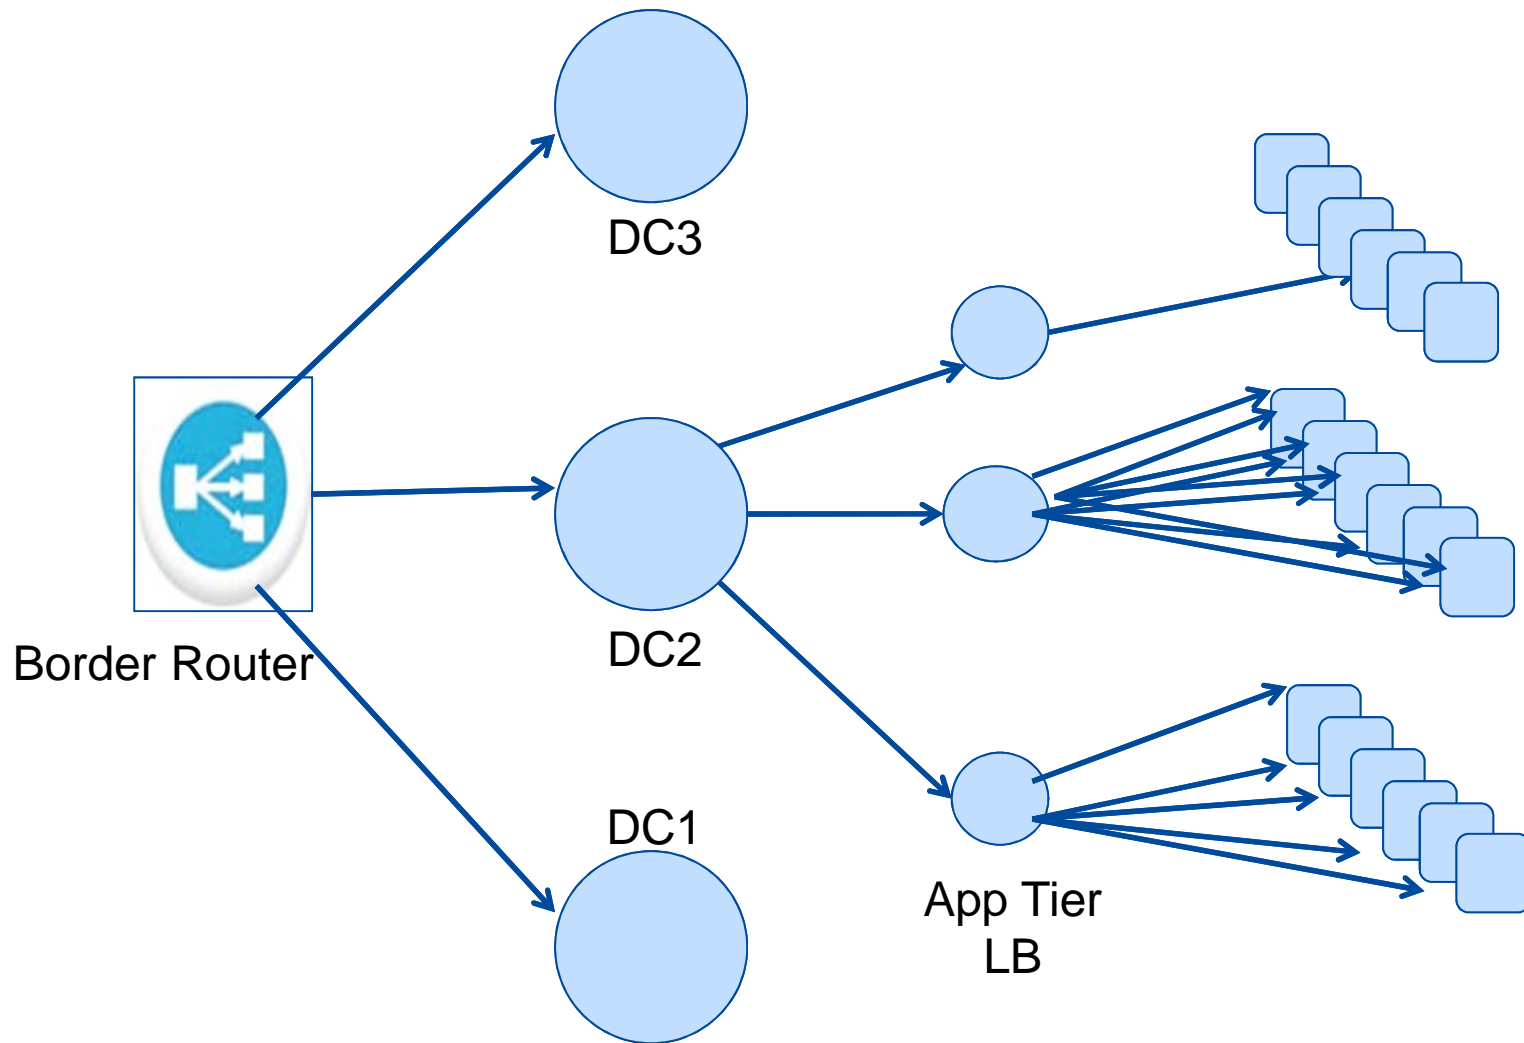
Why Cloud?

- Infrastructure and platform lifecycle management
- Ease of Operations
 - Monitoring and Remediation
 - Metrics and Analysis
- Capacity management
 - Improved resource utilization
 - Quick Flex up/down

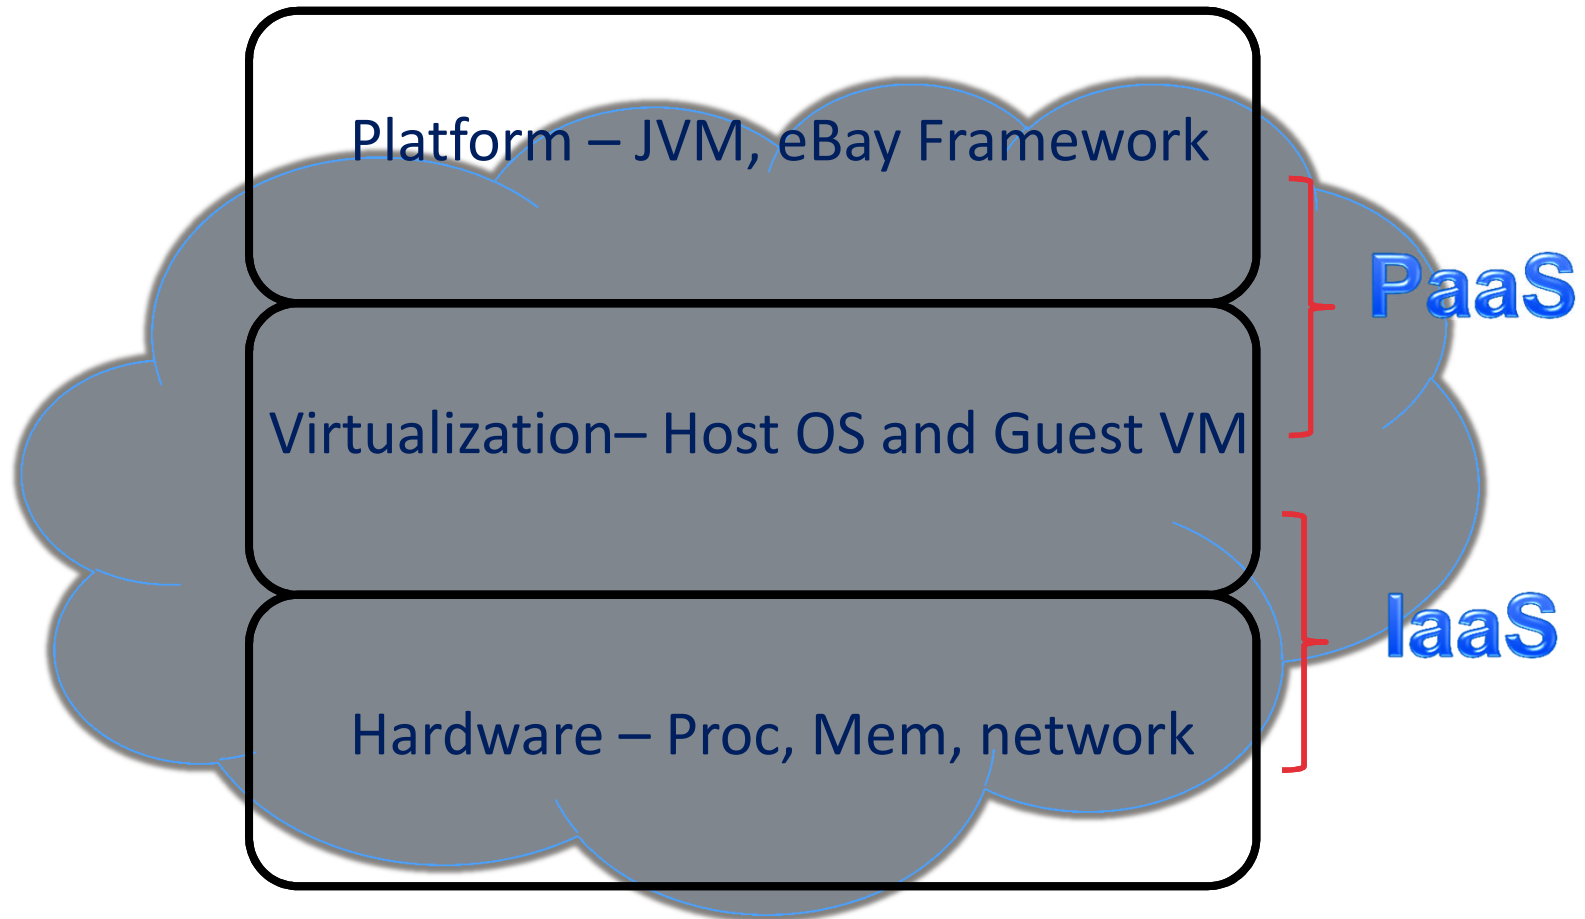
Topology



Routing in data center

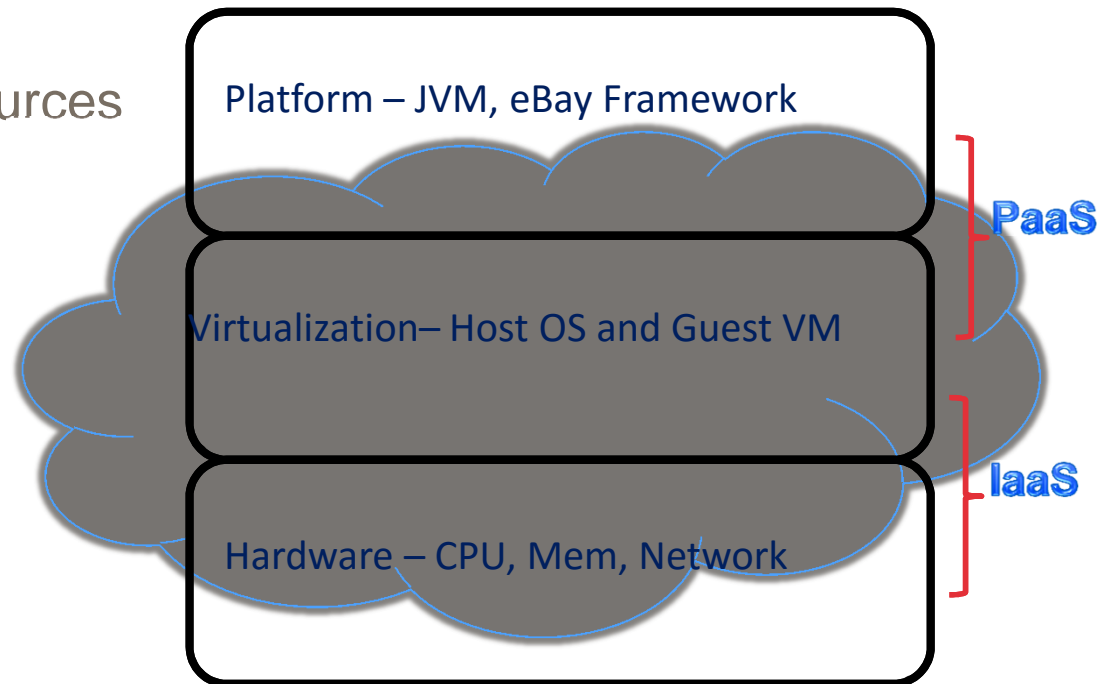


Layers



Challenges

- Heterogeneous hardware
- Virtualization hazards
 - Noisy neighbors
 - Sharing/Stealing of resources
- Platform Integration
 - No Standard metrics
 - Need to be normalized
- Capacity nightmare
 - Scalability
 - Predictability



Traditional Capacity Planning

- Metrics from lab test
- Sometimes from live production
 - Data available only for the operating range
- Extrapolate this for future anticipated traffic

- Doesn't take scalability into consideration
- Potential under-provisioning leads to high transaction response times

Canary - 1

- Red Canary

- Always DR level traffic
- Eventually will lead to right size

- Fixed size hardware in cloud environment*
- Inevitable wasted capacity*
- Utilization profile of canary vs. regular machine

- AutoStress tool – characterize the application

Canary - 2

- Brown Canary

- Application characterization
- Auto-stress tool to scale this step of process
- More accurate estimate of Max throughput

- *Application characteristics change over time*
- *Inevitable wasted capacity*

Utilization profile

- Utilization percentage on VM
- Relative to capacity
- Estimate of percentage of capacity wasted

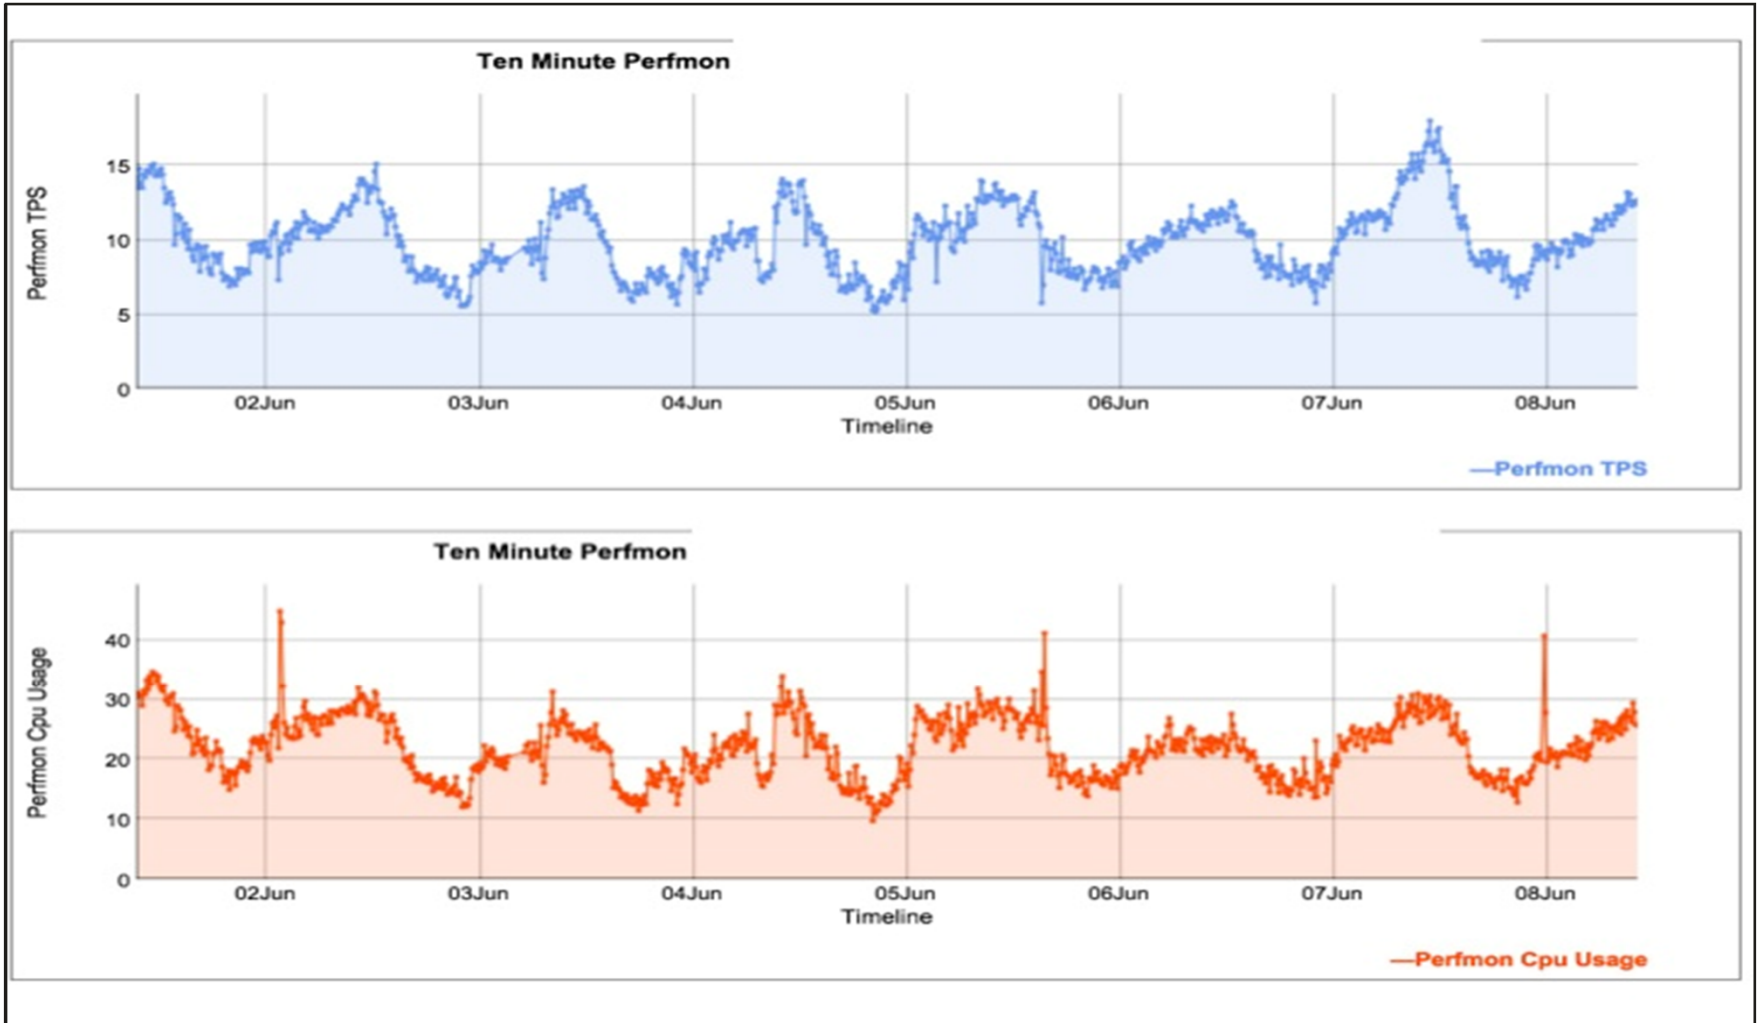
Two Alternatives

- Pre-determine number of machines needed and send traffic to all machines equally.
 - All machines will be almost equally utilized.
 - Hard to tell when to add a new machine.
 - Machines may be heterogeneous.
- Send traffic to machines in an ordered way.
 - Send traffic to machine 1, send to machine 2 when 1 saturates, etc.
 - Unequal usage but when traffic reaches last machine, add capacity.

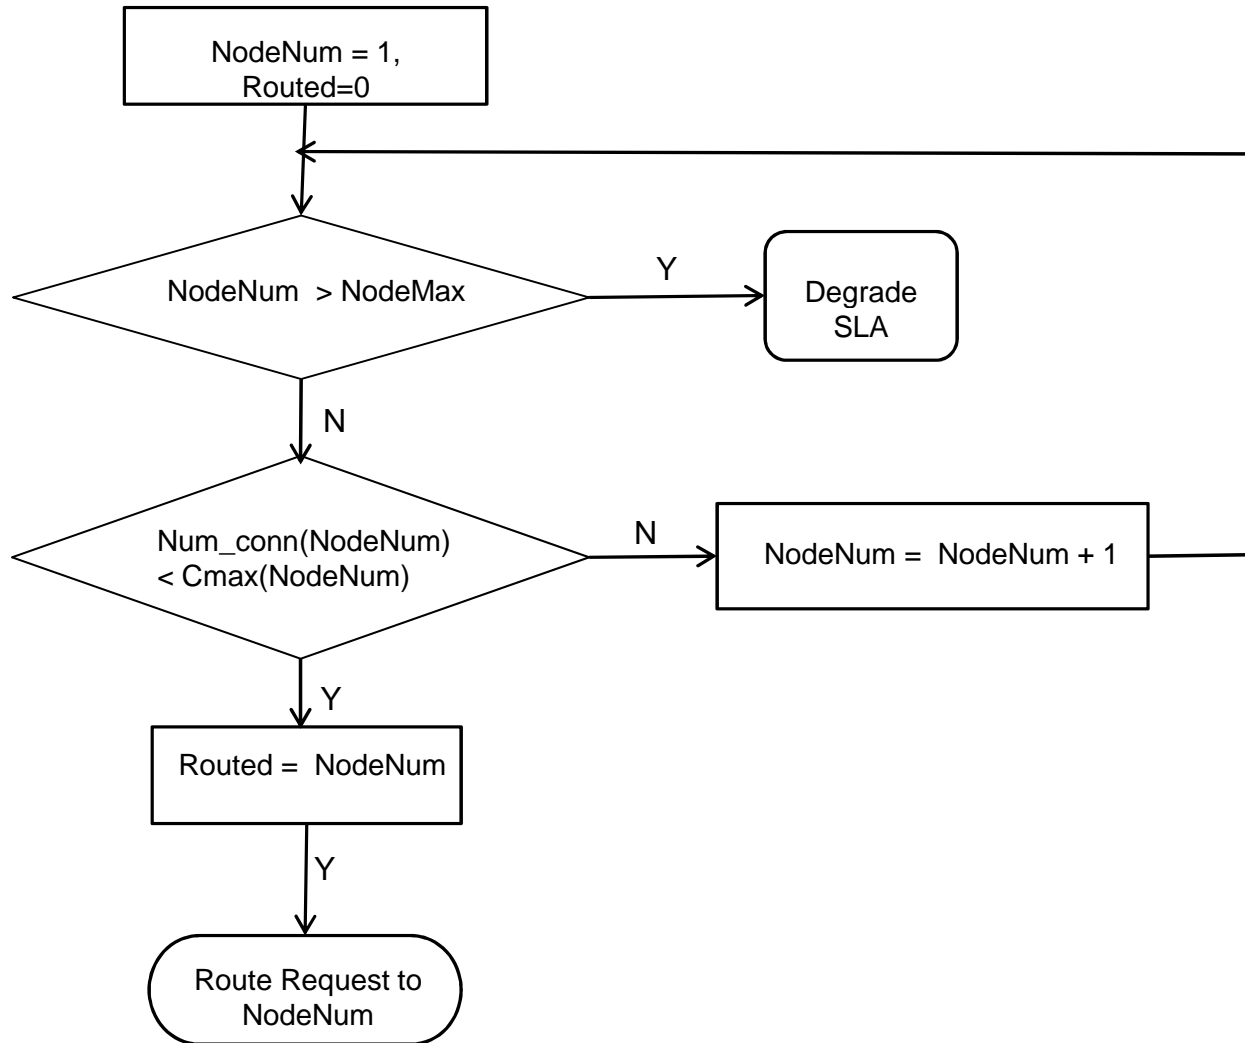
New Idea

- Send all traffic to minimum number of machines.
- Use response time metric as feedback.
- When response time worsens, add machines.


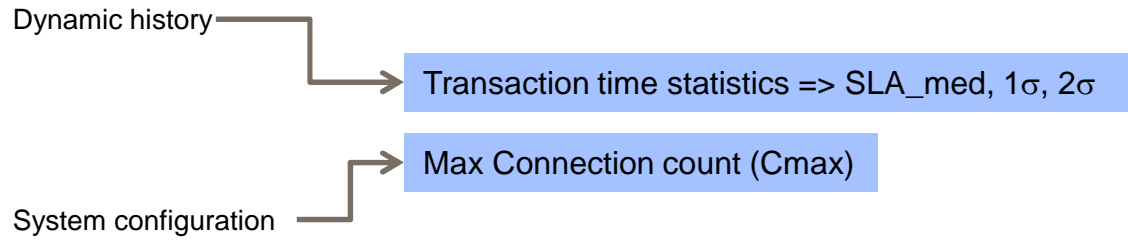
Typical Traffic Pattern



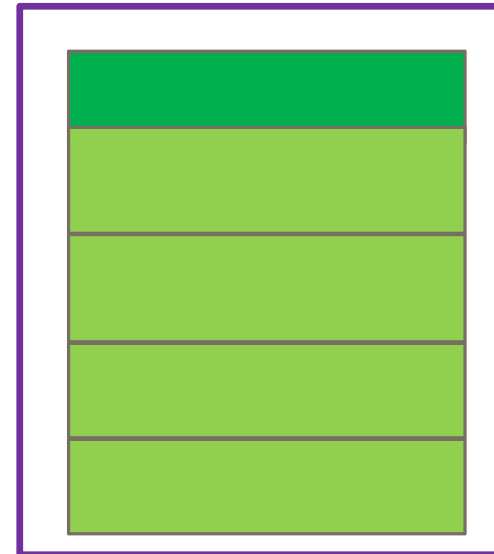
How the flow works – Node Selection



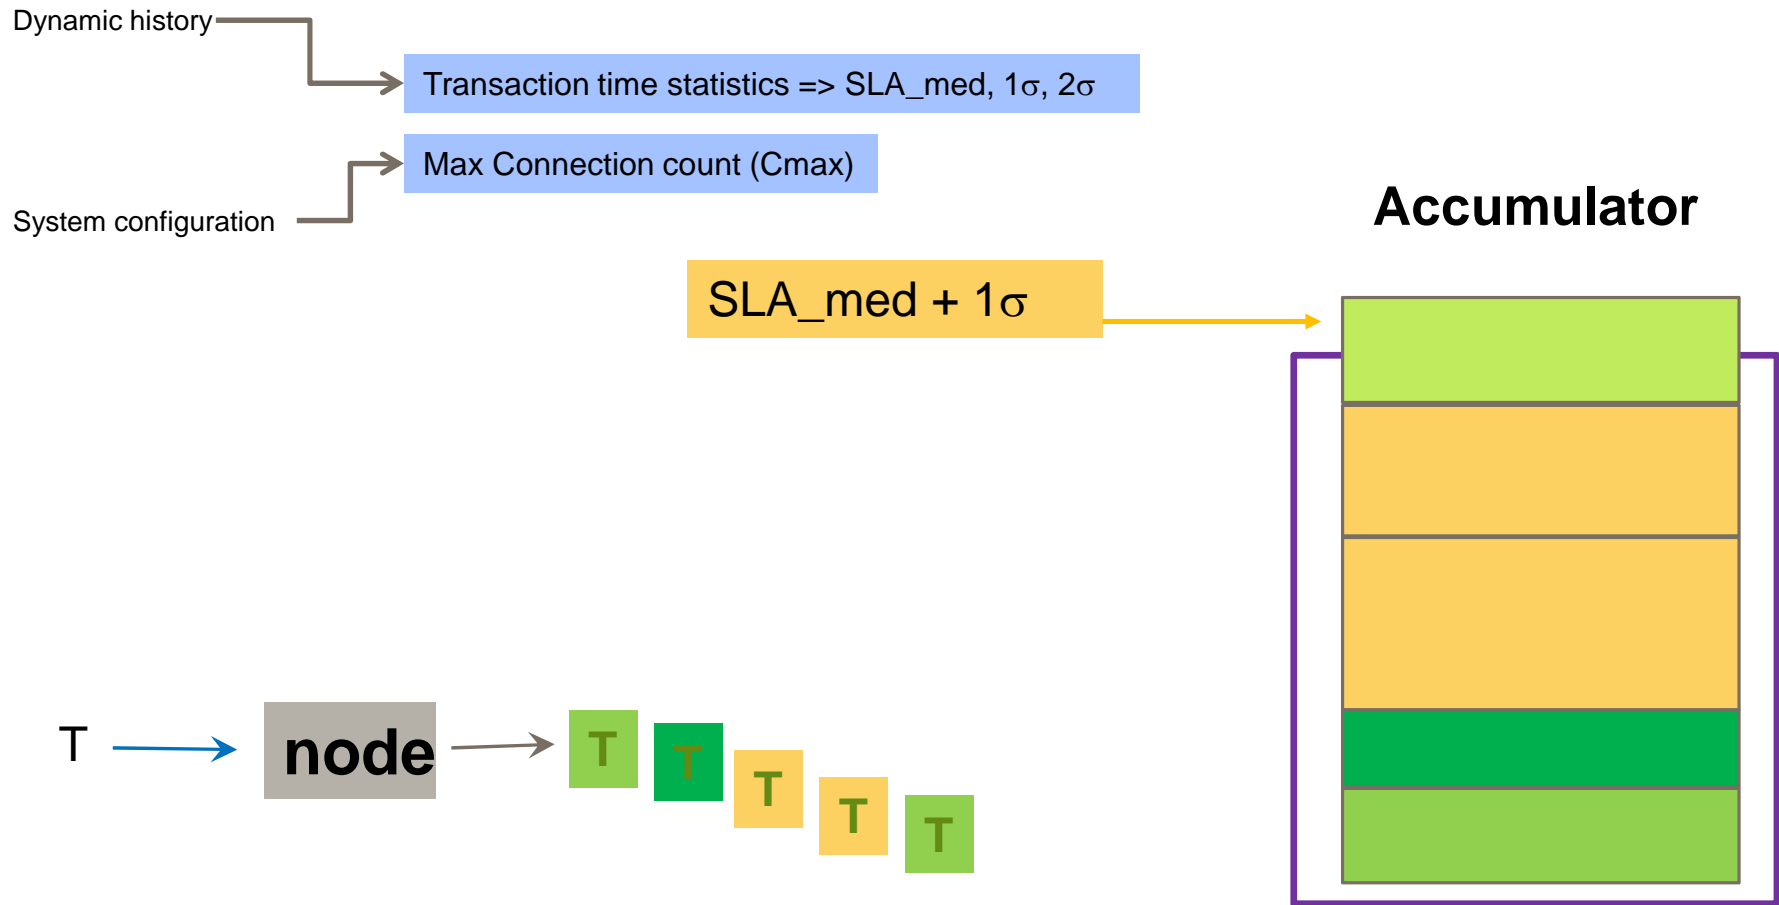
Core of the algorithm



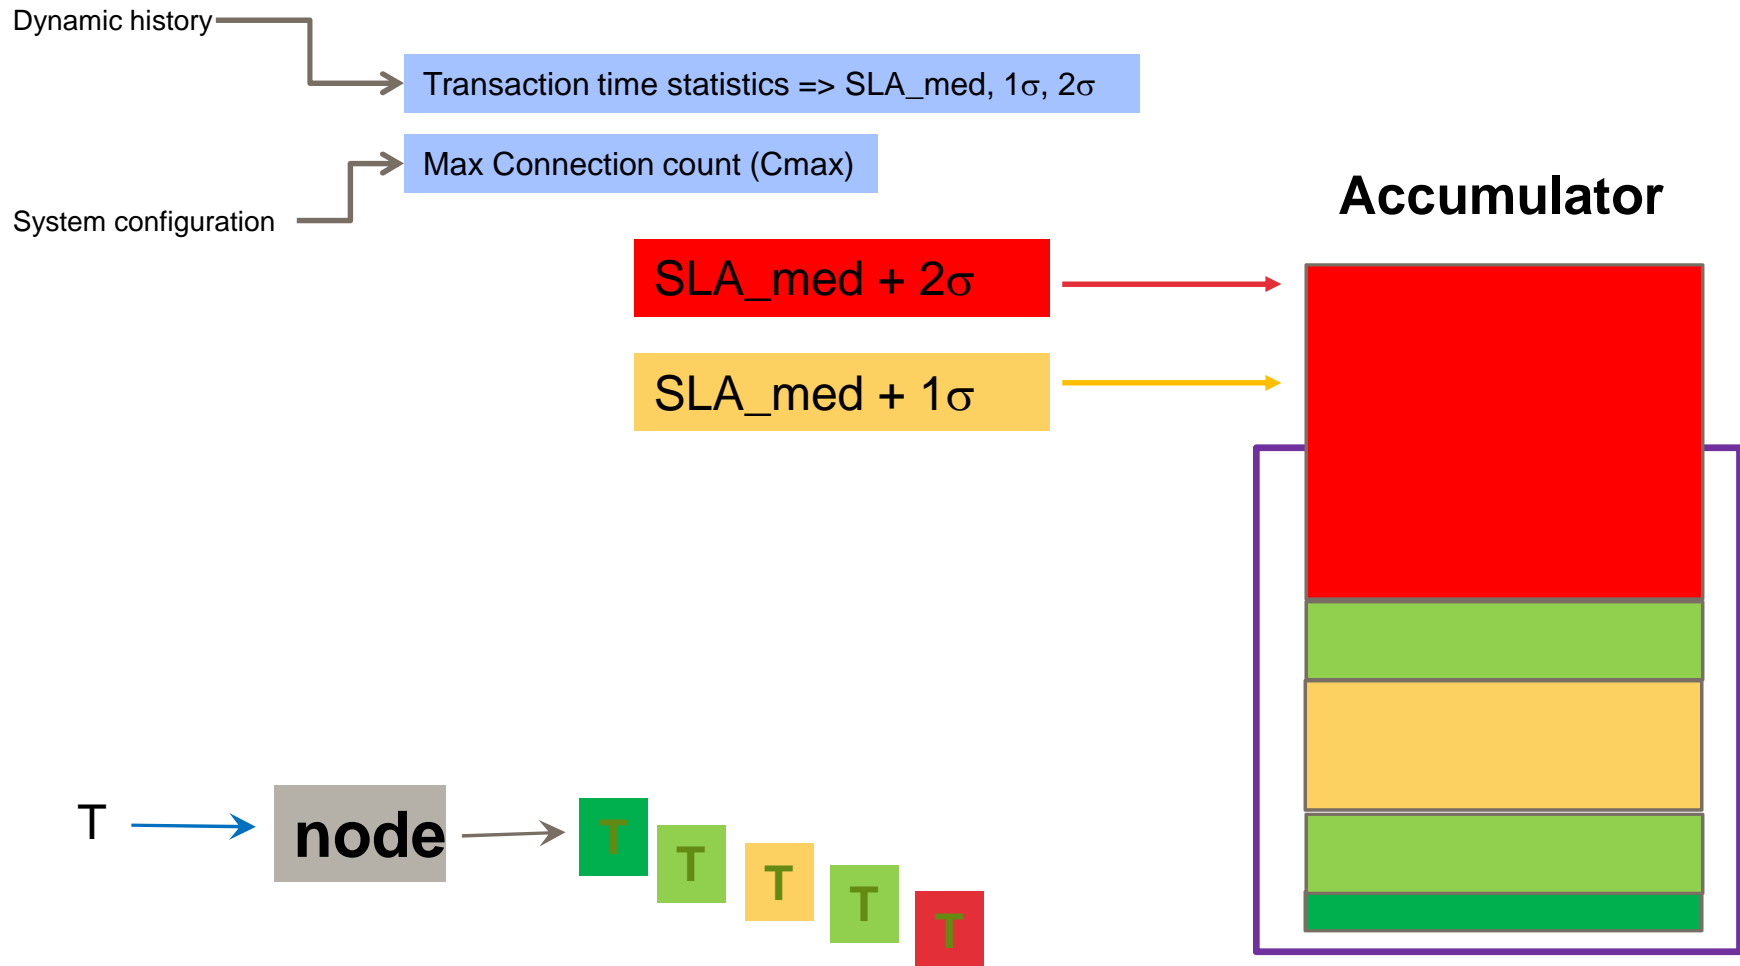
Accumulator



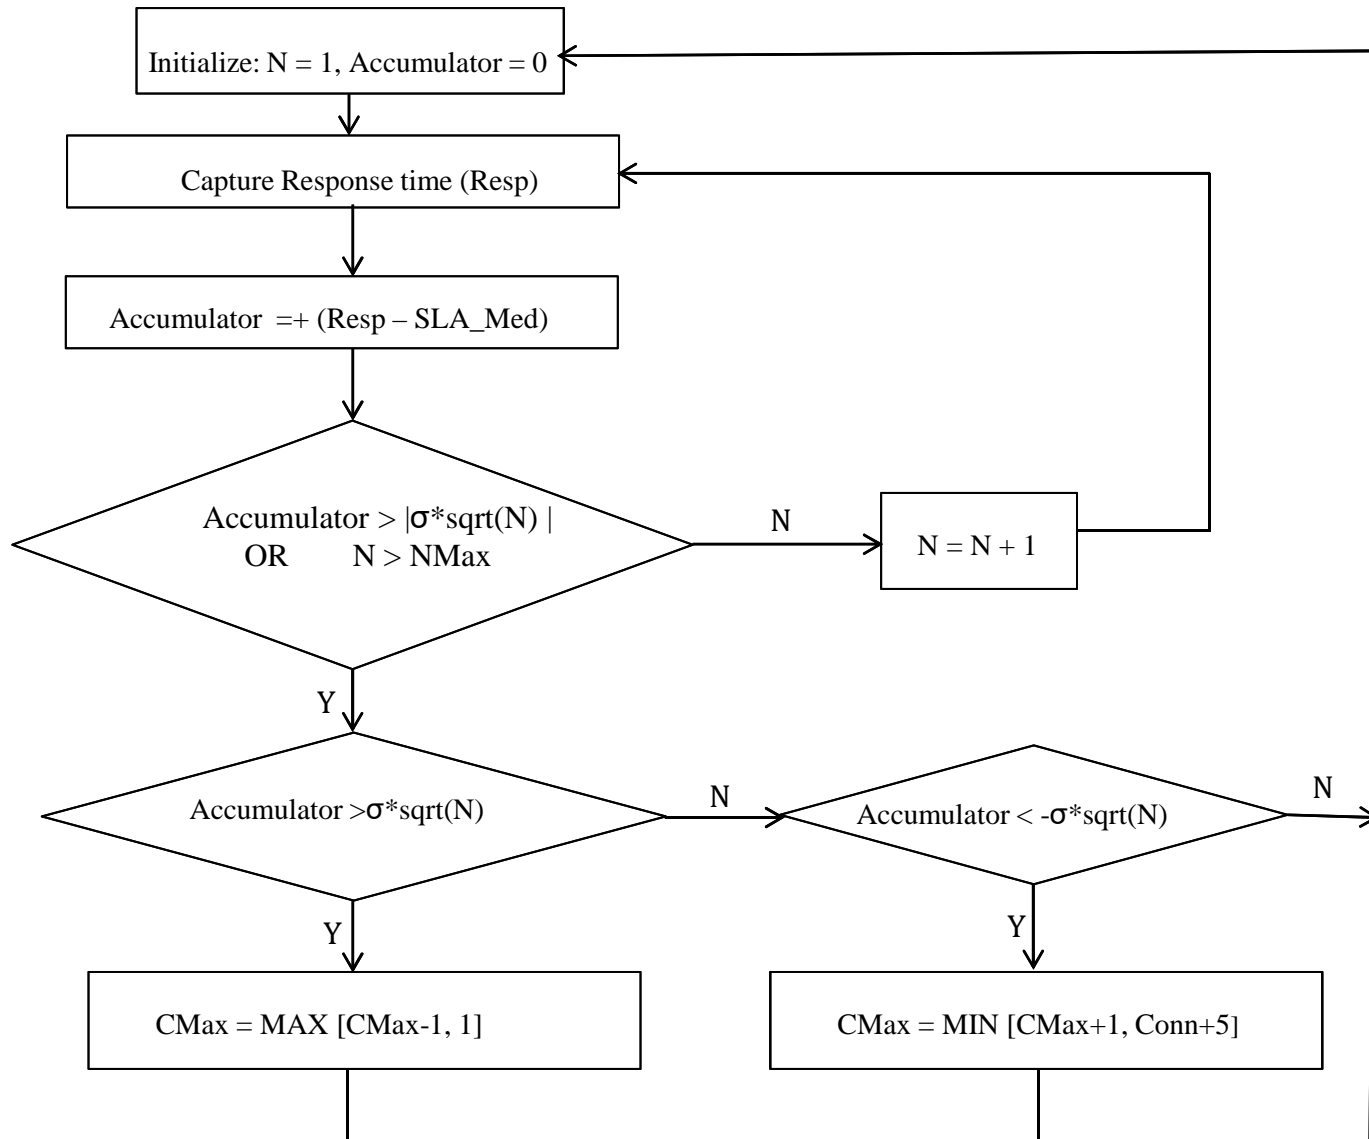
Core of the algorithm



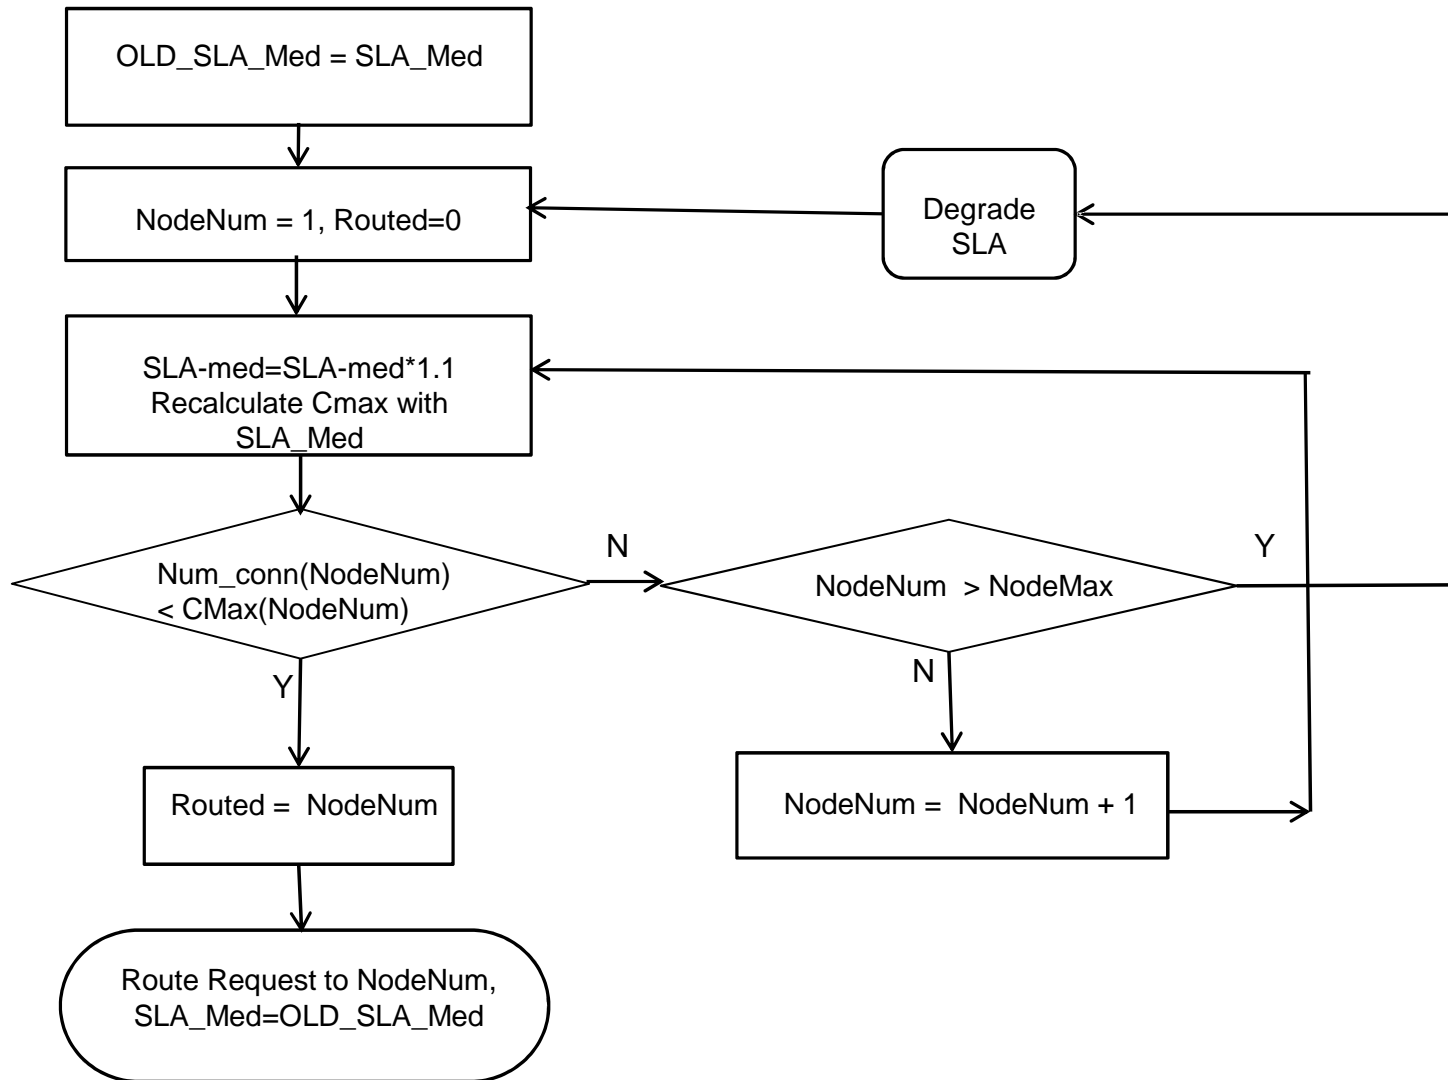
Core of the algorithm



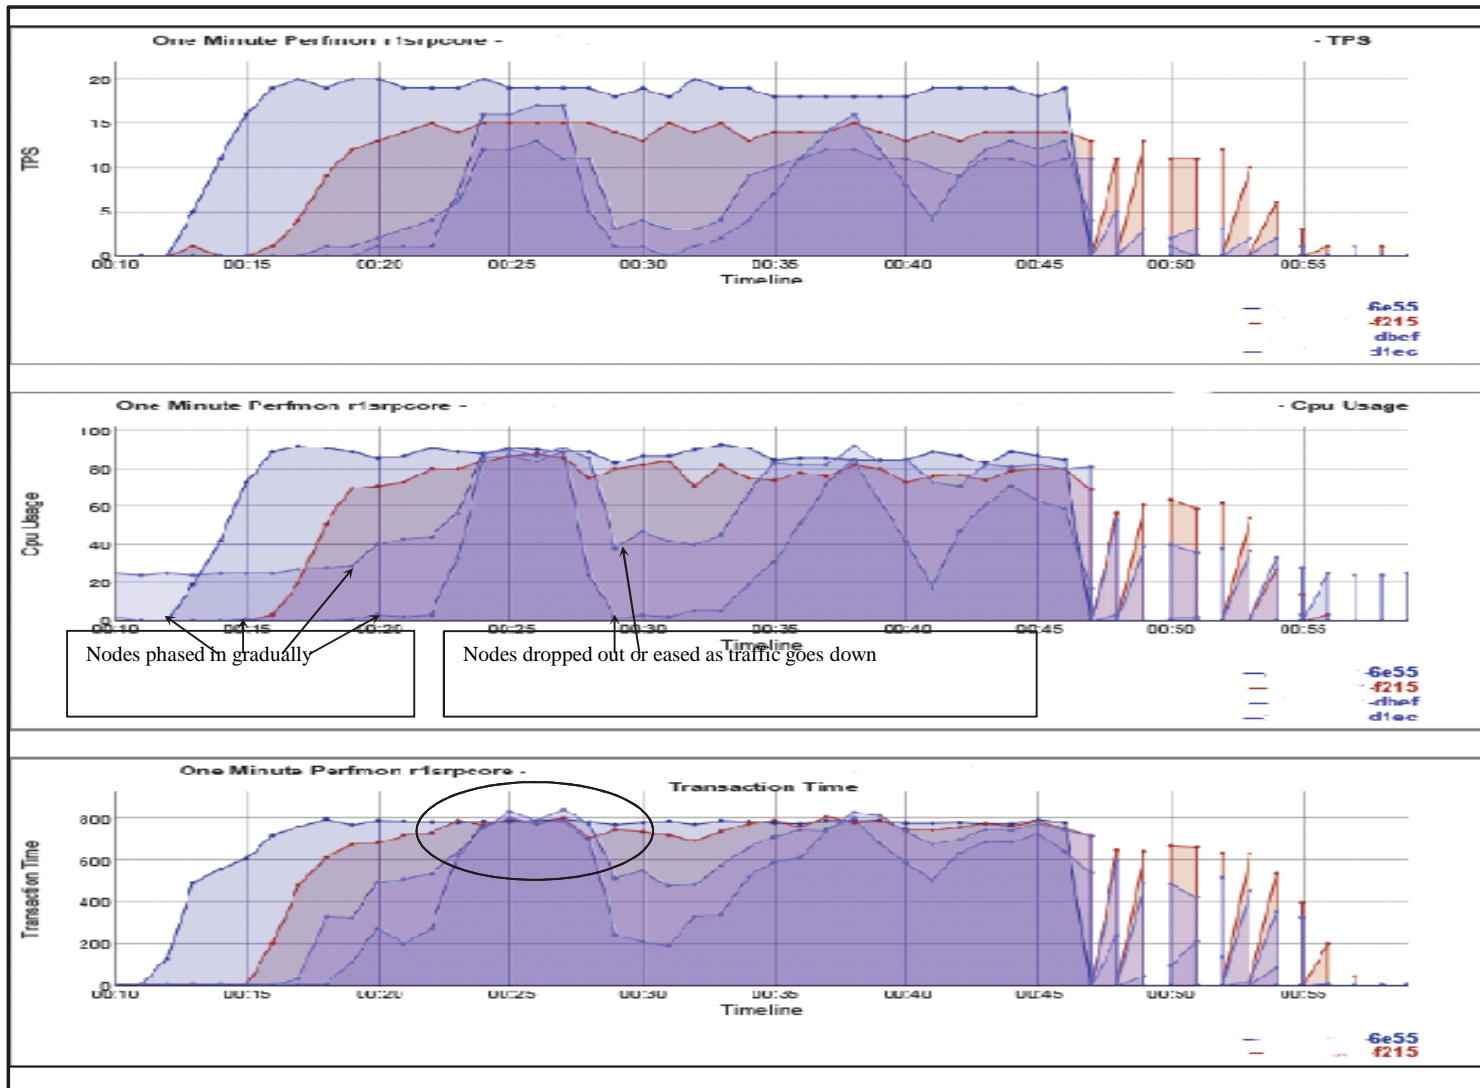
How the flow works – at each node




Degraded Operation to accommodate bursts



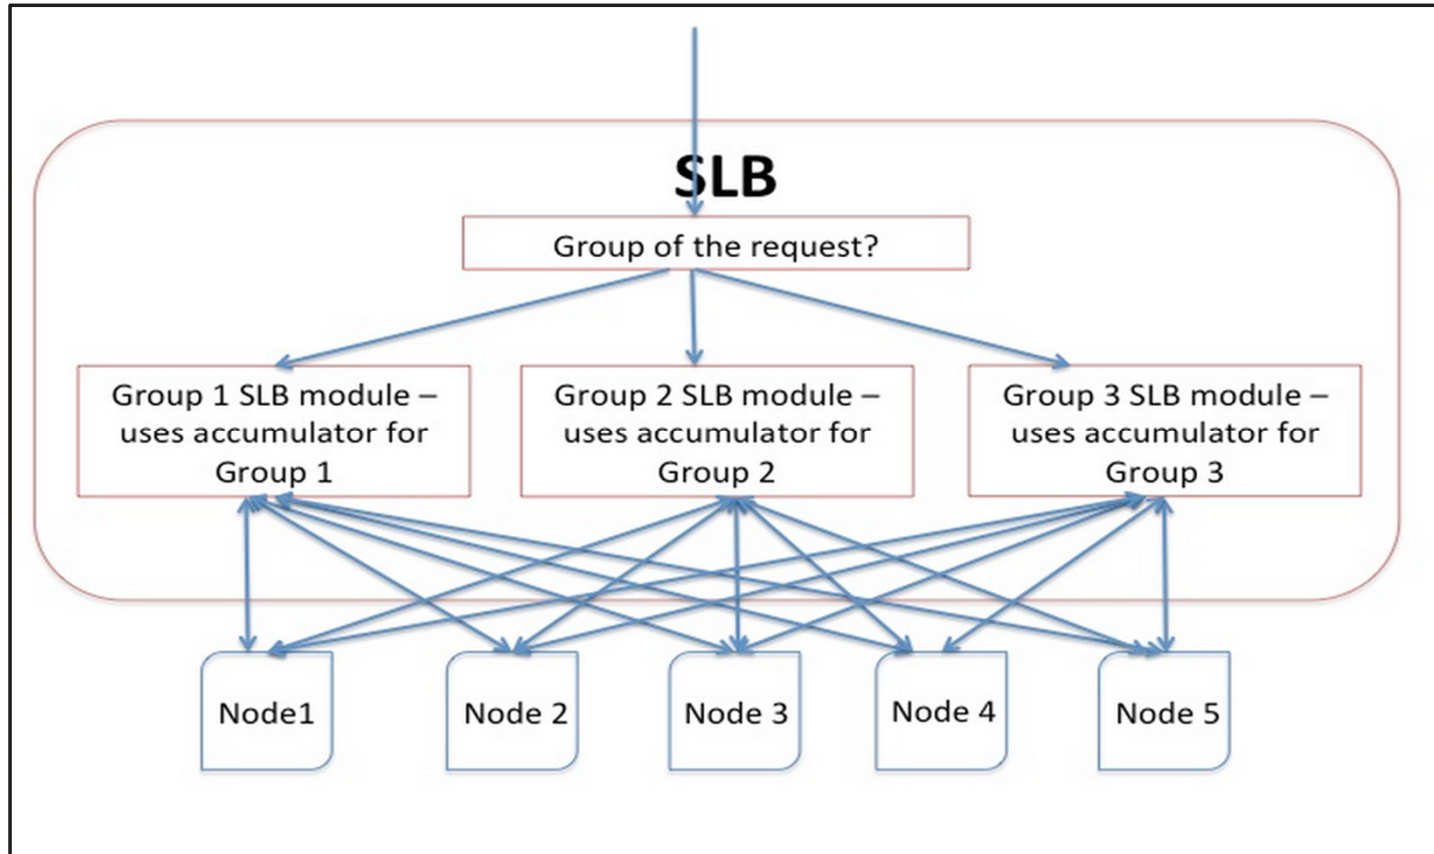
Well behaved load



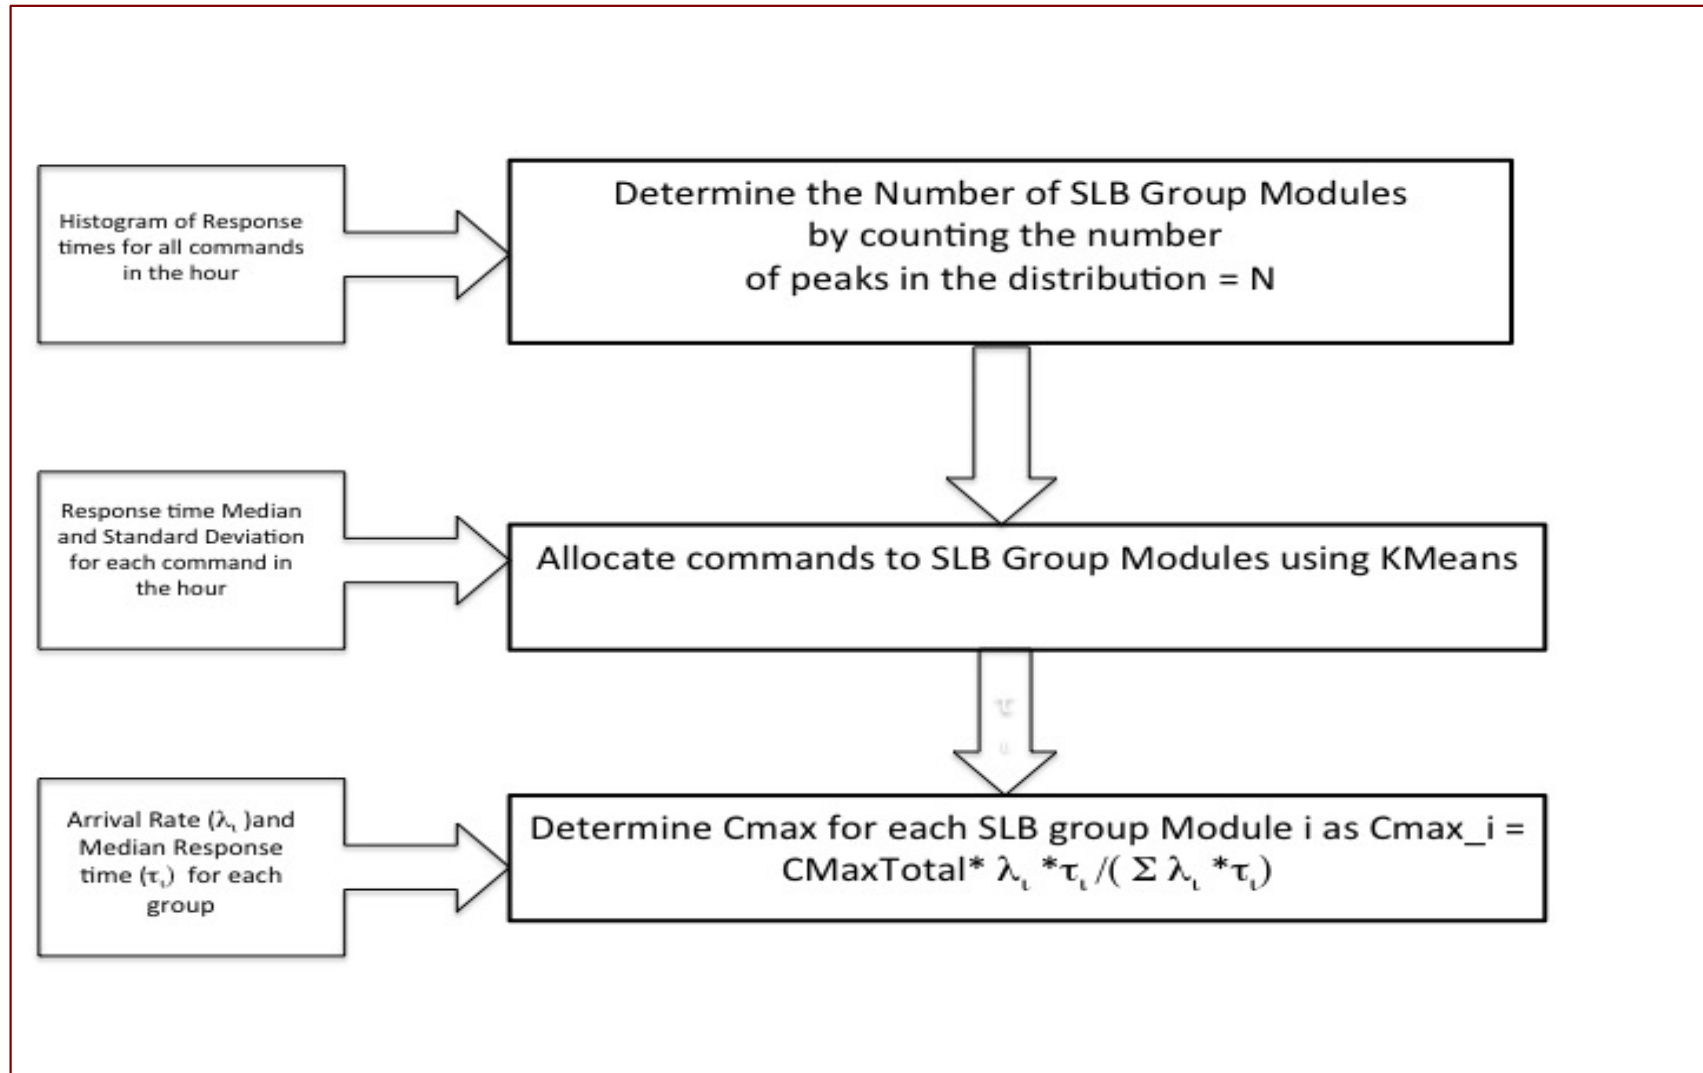
Inadequate capacity in-time – Degraded response time



Extension to Heterogeneous environment



How to create groups?




Comparisons by command type

Command	Median			P95			Count
	Target	Ref	T-R in %	Target	Ref	T-R in %	Ref
AdvS	123	109	12.84	262	238	10.08	3892
AllD	103	112	-8.04	208	218	-4.61	5540
ChsM	457	459	-0.44	1228	1028	19.44	8960
Cust	43	39	10.26	138	138	0.00	1206
FavS	19	222	-91.44	106	417	-74.47	1010
FndH	637	613	3.92	2236	2881	-22.37	43
FndM	90	96	-6.25	1098	5315	-79.34	593
GetC	76	75	1.33	76	75	1.33	1
JsDi	569	559	1.79	1054	1066	-1.11	11625
Prev	163	160	1.88	260	242	7.44	722
RecC	363	368	-1.36	517	544	-4.96	82951
SvSD	320	322	-0.62	549	558	-1.61	273873
SePr	477	477	0.00	682	686	-0.45	1631
SRPR	679	676	0.44	1400	1432	-2.23	1444307
SRSS	560	561	-0.18	1111	1170	-5.04	200592
SelO	585	579	1.04	1060	1059	0.09	62547
SimI	575	583	-1.37	915	985	-7.11	131498
V4Aj	244	104	134.62	323	196	64.96	4
Vero	832	750.5	10.86	2830	907	212.14	2
ZipP	44	42	4.76	96	71	35.21	426
TOTAL							2231423

Thank you!!

**Now,
Open for Discussion**



The power of our platform

Our technology transformation enables us to:

- Process payments at tremendous scale
- Accelerate the innovation of new products
- Engage world-class developers & technologists