Satellite at Scale

How Walmart manages the massive IT infrastructure at the heart of the retail giant

Darin Lively, Brian Ameling, & John Henley
2018-05-08

Any reference in this presentation to any specific commercial product, process, or service, or the use of any trade, firm or corporation name is for information and convenience purposes only, and does not constitute an endorsement or recommendation by Walmart Inc.

PLEASE NOTE

This presentation contains specific configuration and parameters optimized to Walmart's environment. These are likely not the exact configurations or parameters that you will use in your environment. But if they can scale, we know you can, too.

WALMART

SCALE

Walmart's Satellite 6 Journey

PART ONE	**Walmart & Red Hat Joint Engineering**
PART TWO	**Large-Scale Satellite Operations**
PART THREE	**Satellite Infrastructure Considerations**
PART FOUR	**Managing Satellite Clients at Scale**
PART FIVE	**Operational Improvements**

redhat

Walmart

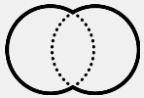

JOINT ENGINEERING EFFORT

COLLABORATIVE PARTNERSHIP

REPEATABLE PROCESS, OPEN TO ANYONE

IMPROVING SATELLITE FOR ALL

Large-Scale Satellite Operations

WHAT WALMART BUSINESS SCALE LOOKS LIKE

2.3M
ASSOCIATES

28
COUNTRIES

270M
CUSTOMERS/WEEK

11,700
STORES

150+
DISTRIBUTION CENTERS

WHAT WALMART SATELLITE SCALE LOOKS LIKE

60K
HOSTS

64
CONTENT VIEWS

9
LIFECYCLE ENVIRONMENTS

3.9K
PULP REPOS

234K
PACKAGES

182K
PACKAGE DOWNLOADS/ HOUR PEAK

46
ACTIVATION KEYS

22
CAPSULES

1
SATELLITE

Puppet Client Total

93393

Puppet Clients Unchanged

RHEL Subscriptions Consumed

Satellite Client Total

60019

Puppet Clients

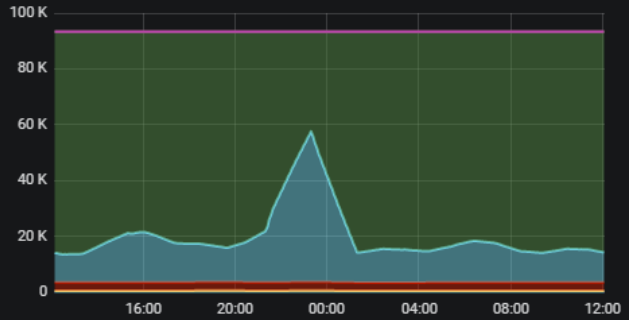

	max	current
NOOP	0	0
Unreported	23	16
Unresponsive	641	443
Failed	3.0 K	3.0 K
Changed	53.9 K	10.8 K
Unchanged	79.8 K	79.2 K
Total	93.4 K	93.4 K

Satellite Clients & Provisioning

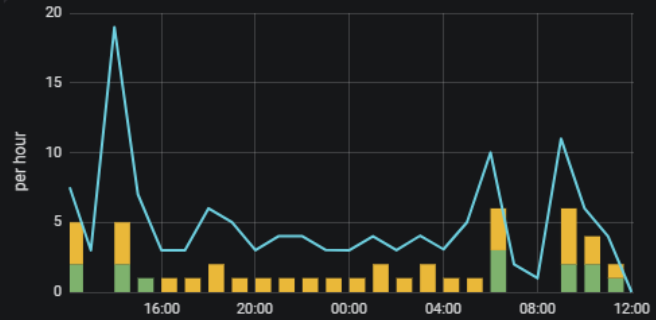

	max	total
Builds (Sat6)	3	13
Builds (SDI-Portal)	4	33
Registrations	14	78

Puppet Code Deployments

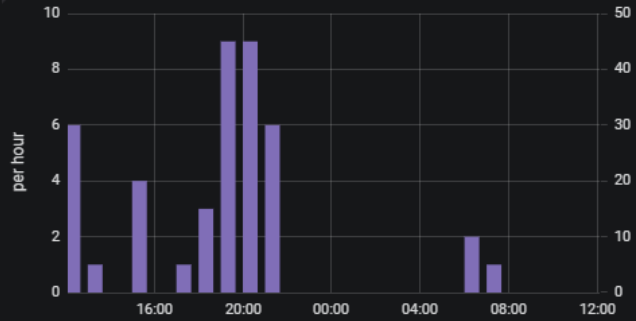

	max	avg
Total/hour	9	2
Cumulative (right-y)	42	30

Satellite Package Downloads

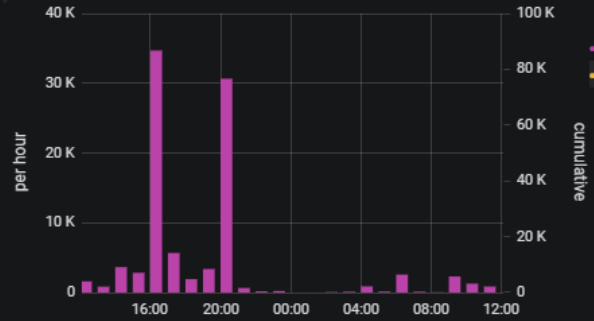

	max	avg
Total/hour	34.7 K	3.8 K
Cumulative (right-y)	94.8 K	67.0 K

Satellite Infrastructure Considerations

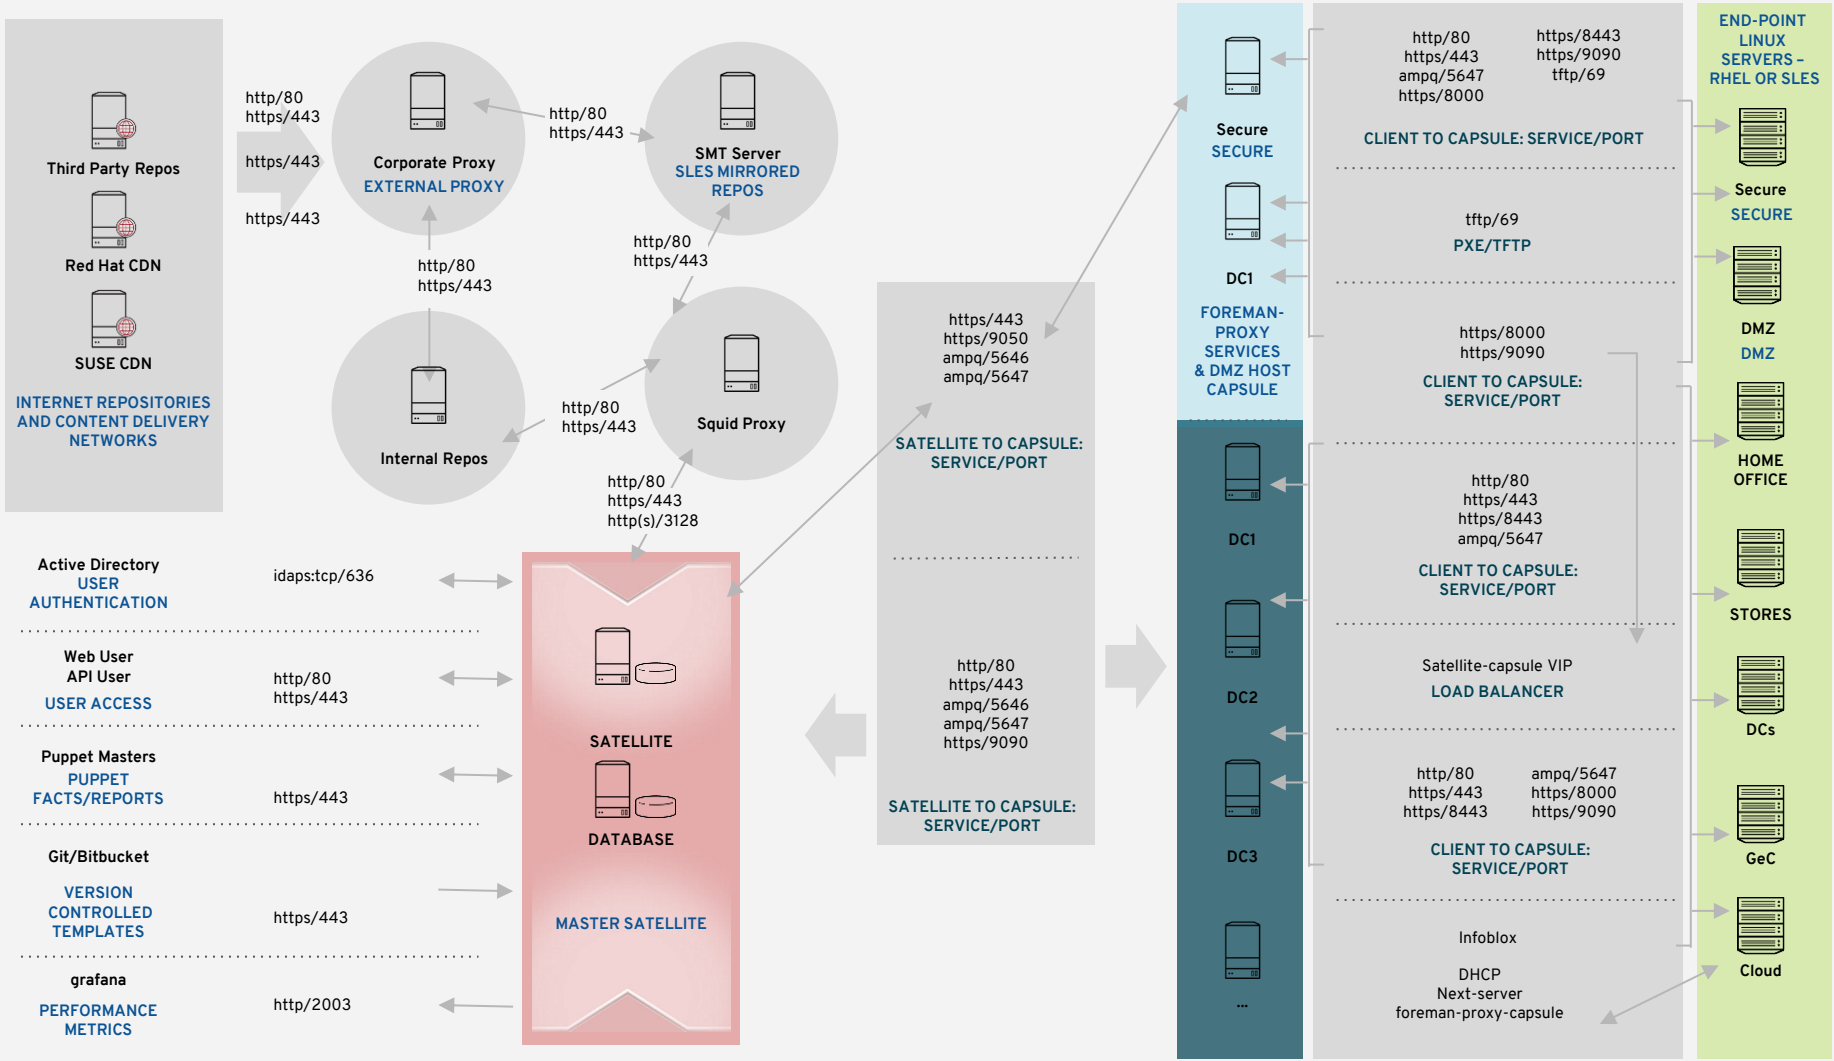
Third Party Repos
Red Hat CDN
SUSE CDN
INTERNET REPOSITORIES AND CONTENT DELIVERY NETWORKS
http/80
https/443
https/443
https/443
Corporate Proxy
EXTERNAL PROXY
http/80
https/443
SMT Server
SLES MIRRORED REPOS
http/80
https/443
http/80
https/443
Internal Repos
http/80
https/443
Squid Proxy
http/80
https/443
http(s)/3128
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE
https/443
https/9050
ampq/5646
ampq/5647
SATELLITE TO CAPSULE: SERVICE/PORT
http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT
Secure
SECURE
DC1
FOREMAN-PROXY SERVICES & DMZ HOST CAPSULE
DC1
DC2
DC3
...
http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT
tftp/69
PXE/TFTP
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT
Satellite-capsule VIP
LOAD BALANCER
http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
Infoblox
DHCP
Next-server
foreman-proxy-capsule
END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

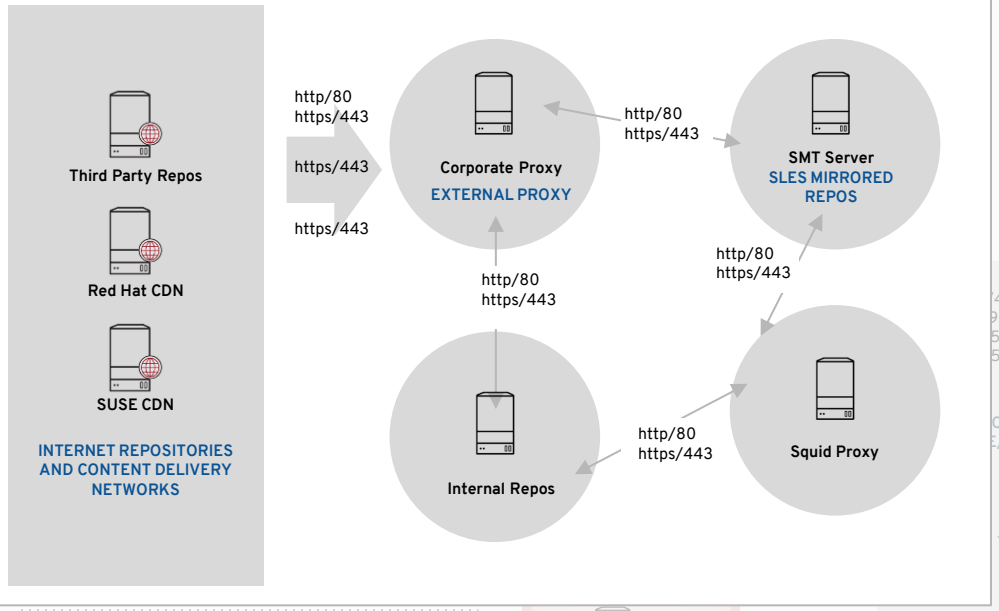

Web User
API User
USER ACCESS
http/80
https/443

Puppet Masters
PUPPET FACTS/REPORTS
https/443

Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443

grafana
PERFORMANCE METRICS
http/2003

SATELLITE
DATABASE
MASTER SATELLITE

http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT

Secure
SECURE
DC1
FOREMAN-PROXY SERVICES & DMZ HOST CAPSULE
DC1
DC2
DC3
...

http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT

tftp/69
PXE/TFTP

https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT

http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT

Satellite-capsule VIP
LOAD BALANCER

http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT

Infoblox
DHCP
Next-server
foreman-proxy-capsule

END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

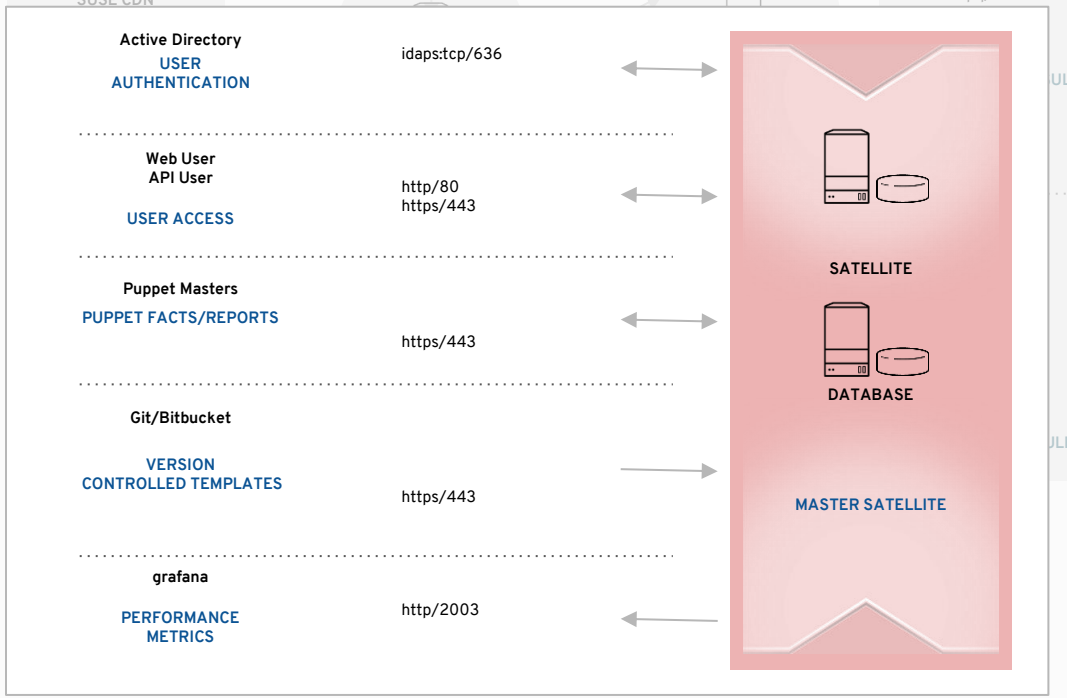
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE

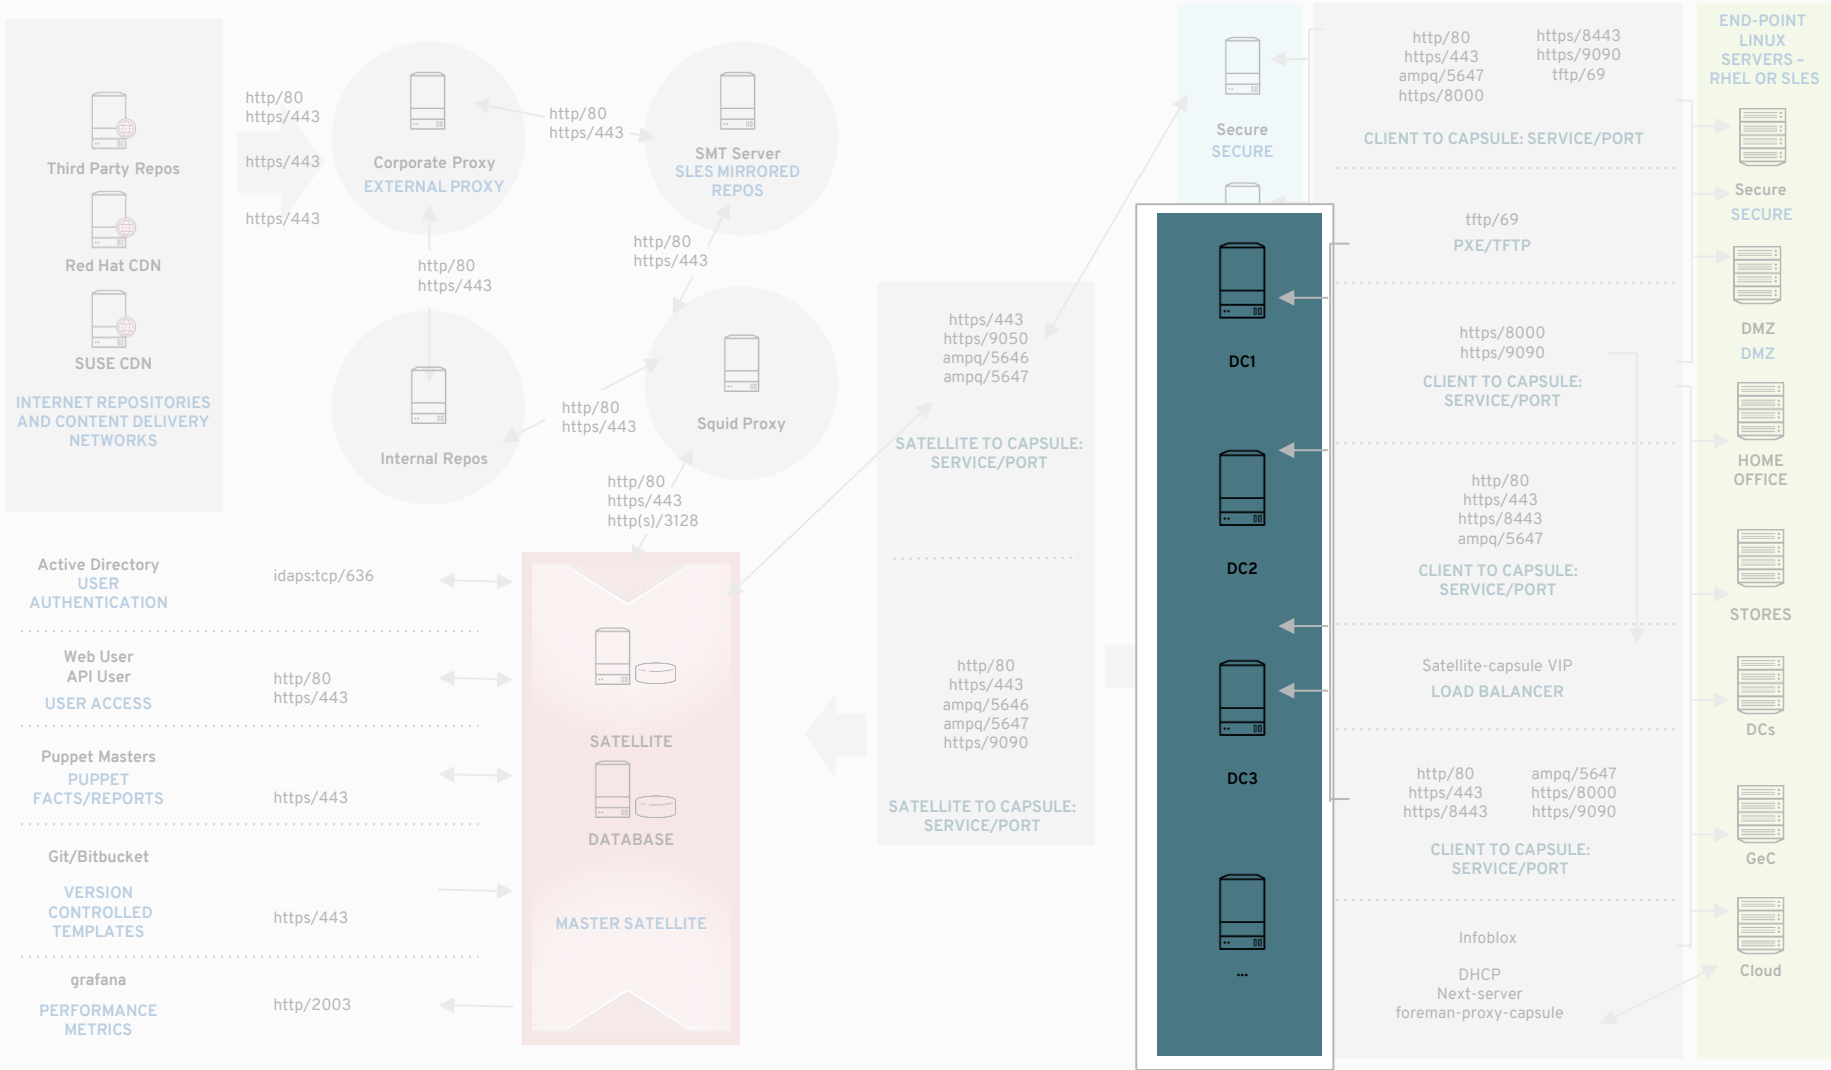
Third Party Repos
Red Hat CDN
SUSE CDN
INTERNET REPOSITORIES AND CONTENT DELIVERY NETWORKS
http/80
https/443
https/443
Corporate Proxy
EXTERNAL PROXY
http/80
https/443
SMT Server
SLES MIRRORED REPOS
http/80
https/443
http/80
https/443
Internal Repos
http/80
https/443
Squid Proxy
http/80
https/443
http(s)/3128
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE
https/443
https/9050
ampq/5646
ampq/5647
SATELLITE TO CAPSULE: SERVICE/PORT
http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT
Secure
SECURE
DC1
DC2
DC3
...
http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT
tftp/69
PXE/TFTP
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT
Satellite-capsule VIP
LOAD BALANCER
http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
Infoblox
DHCP
Next-server
foreman-proxy-capsule
END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

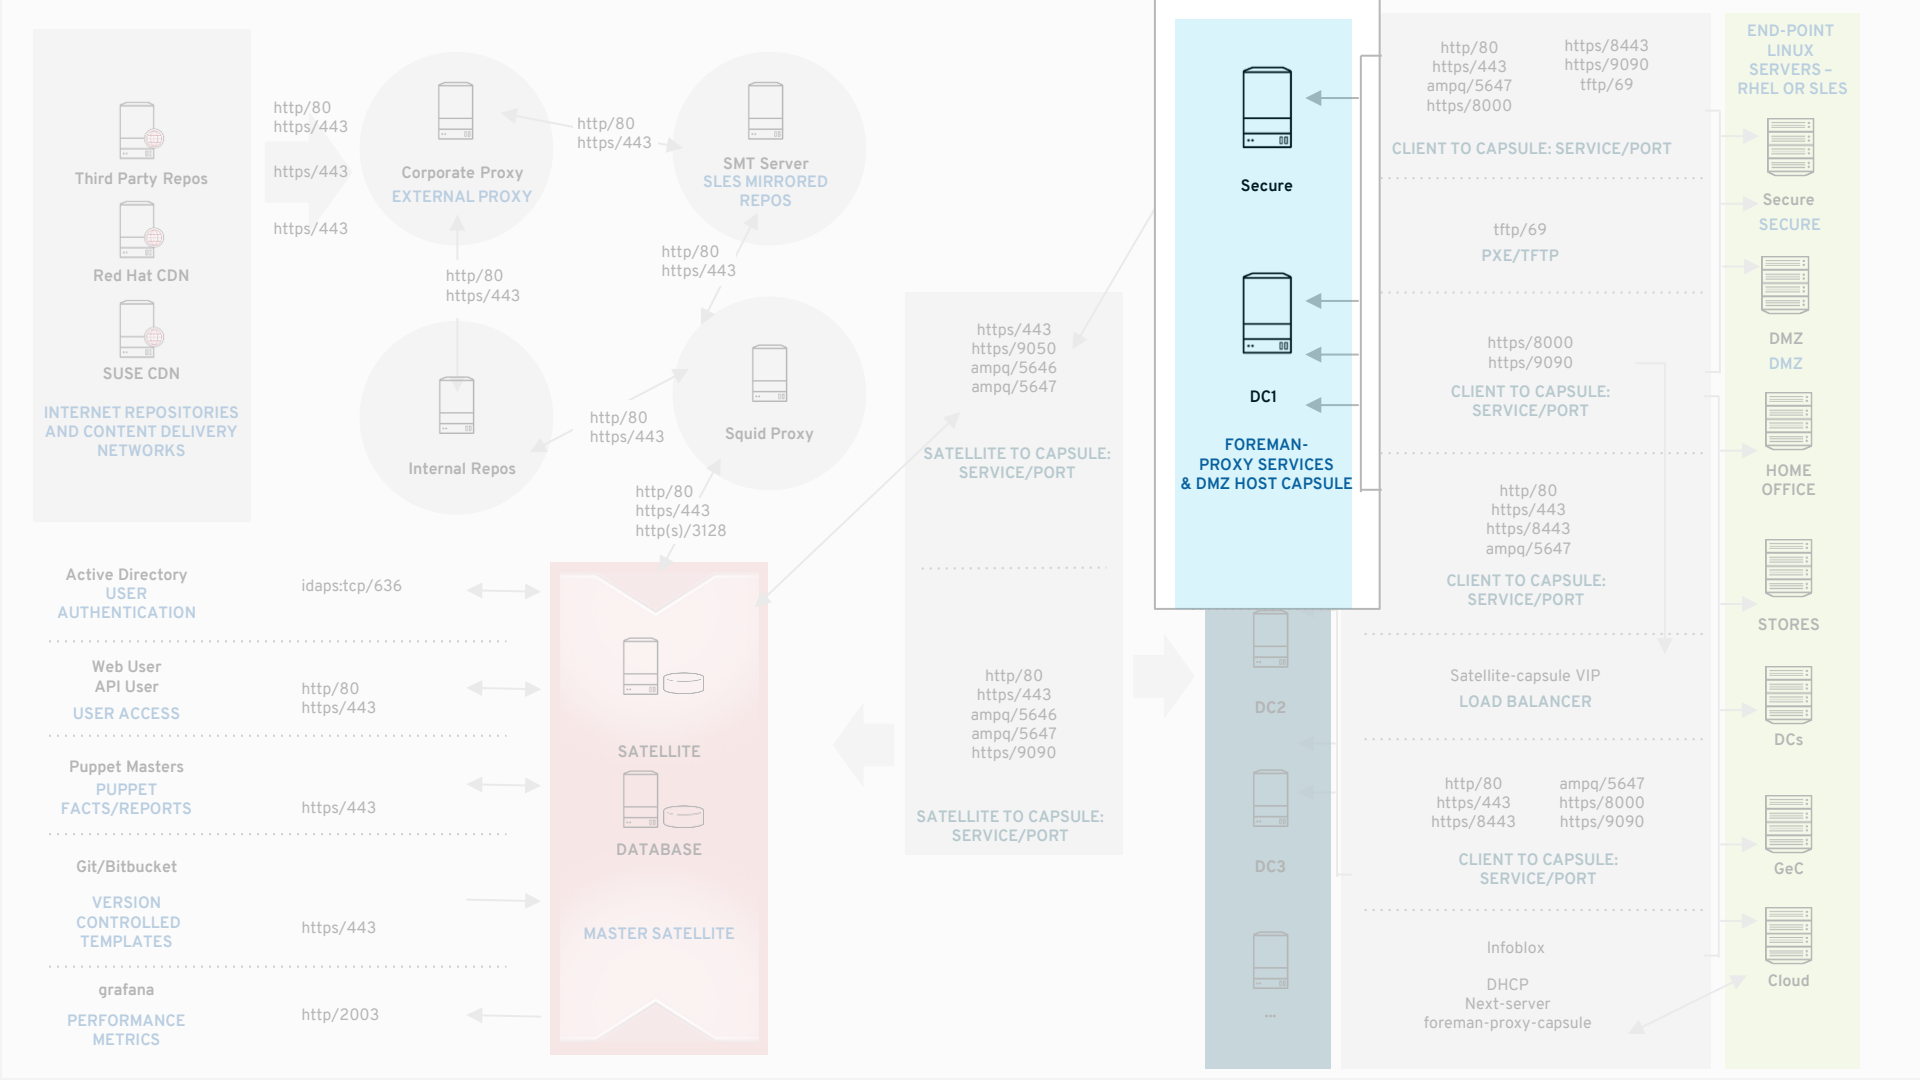
Third Party Repos
Red Hat CDN
SUSE CDN
INTERNET REPOSITORIES AND CONTENT DELIVERY NETWORKS
http/80
https/443
https/443
Corporate Proxy
EXTERNAL PROXY
http/80
https/443
SMT Server
SLES MIRRORED REPOS
http/80
https/443
http/80
https/443
Internal Repos
http/80
https/443
Squid Proxy
http/80
https/443
http(s)/3128
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE
https/443
https/9050
ampq/5646
ampq/5647
SATELLITE TO CAPSULE: SERVICE/PORT
http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT
Secure
DC1
FOREMAN-PROXY SERVICES & DMZ HOST CAPSULE
DC2
DC3
...
http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT
tftp/69
PXE/TFTP
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT
Satellite-capsule VIP
LOAD BALANCER
http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
Infoblox
DHCP
Next-server
foreman-proxy-capsule
END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

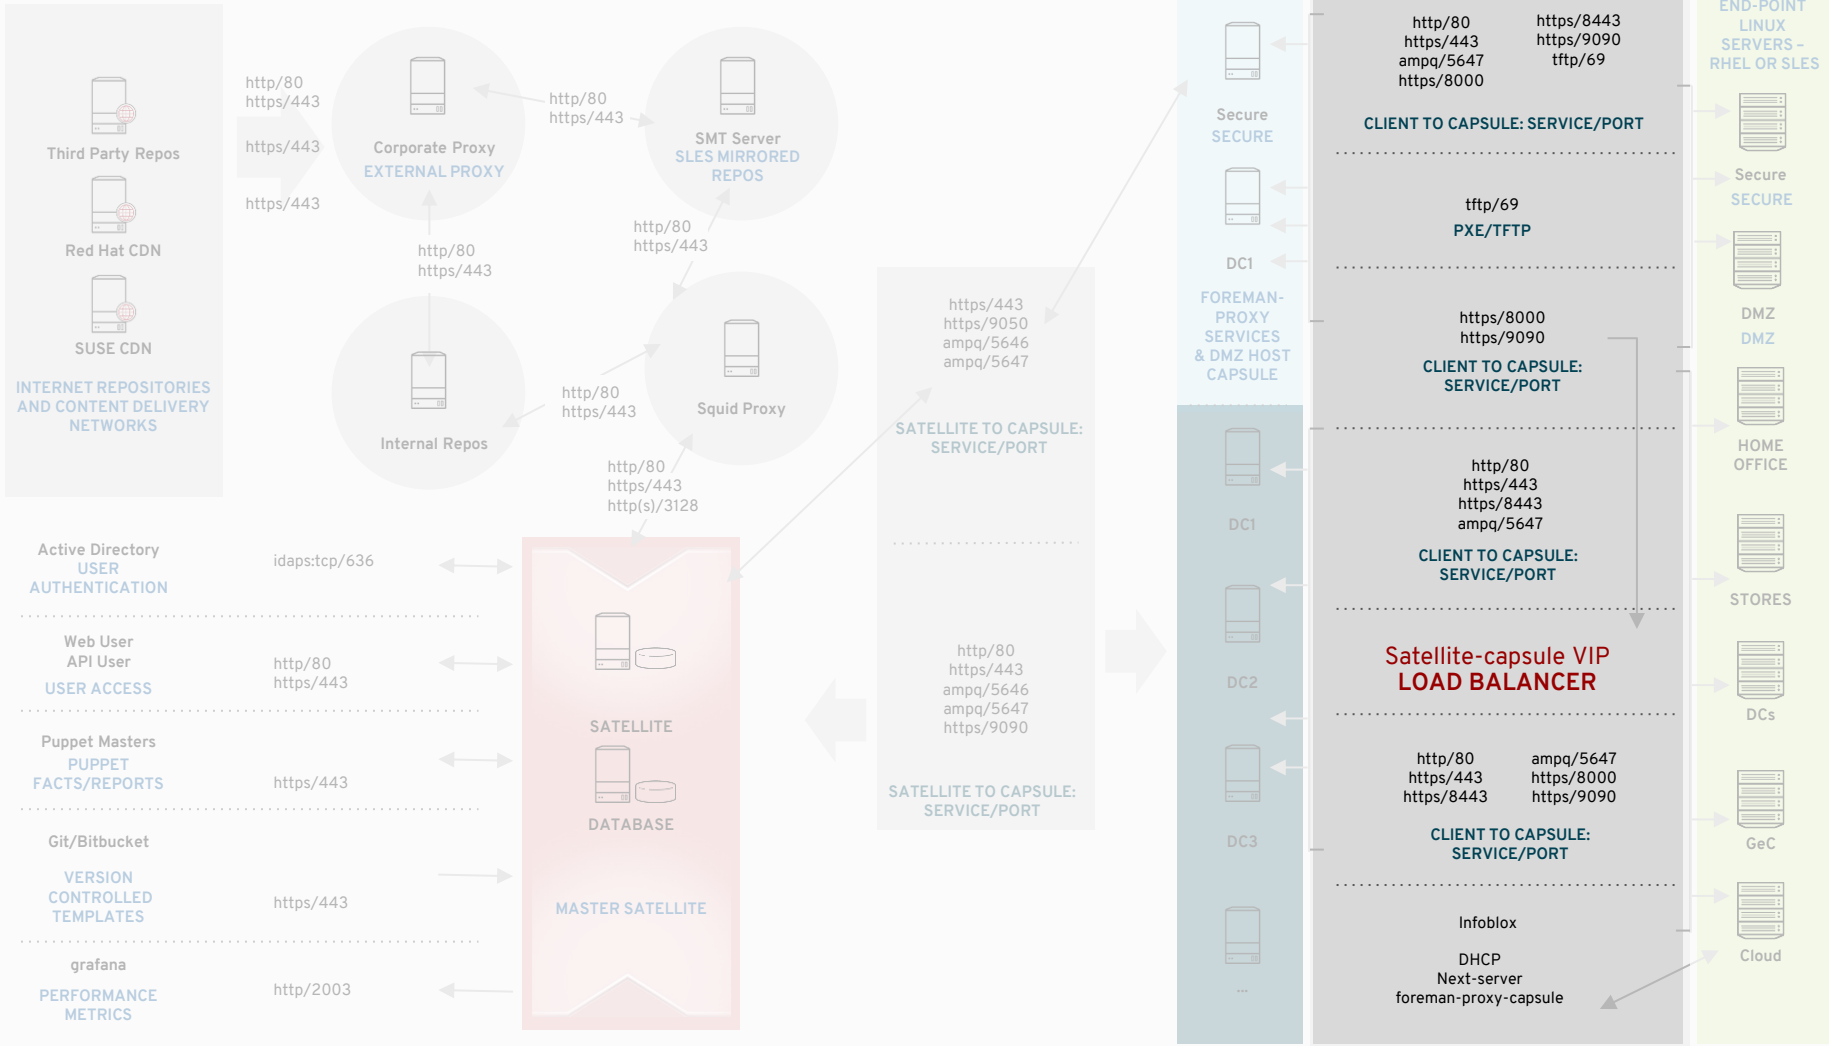
Third Party Repos
Red Hat CDN
SUSE CDN
INTERNET REPOSITORIES AND CONTENT DELIVERY NETWORKS
http/80
https/443
https/443
https/443
Corporate Proxy
EXTERNAL PROXY
http/80
https/443
SMT Server
SLES MIRRORED REPOS
http/80
https/443
http/80
https/443
Internal Repos
http/80
https/443
Squid Proxy
http/80
https/443
http(s)/3128
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE
https/443
https/9050
ampq/5646
ampq/5647
SATELLITE TO CAPSULE: SERVICE/PORT
http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT
Secure
SECURE
DC1
FOREMAN-PROXY SERVICES & DMZ HOST CAPSULE
DC1
DC2
DC3
...
http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT
tftp/69
PXE/TFTP
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT
Satellite-capsule VIP
LOAD BALANCER
http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
Infoblox
DHCP
Next-server
foreman-proxy-capsule
END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

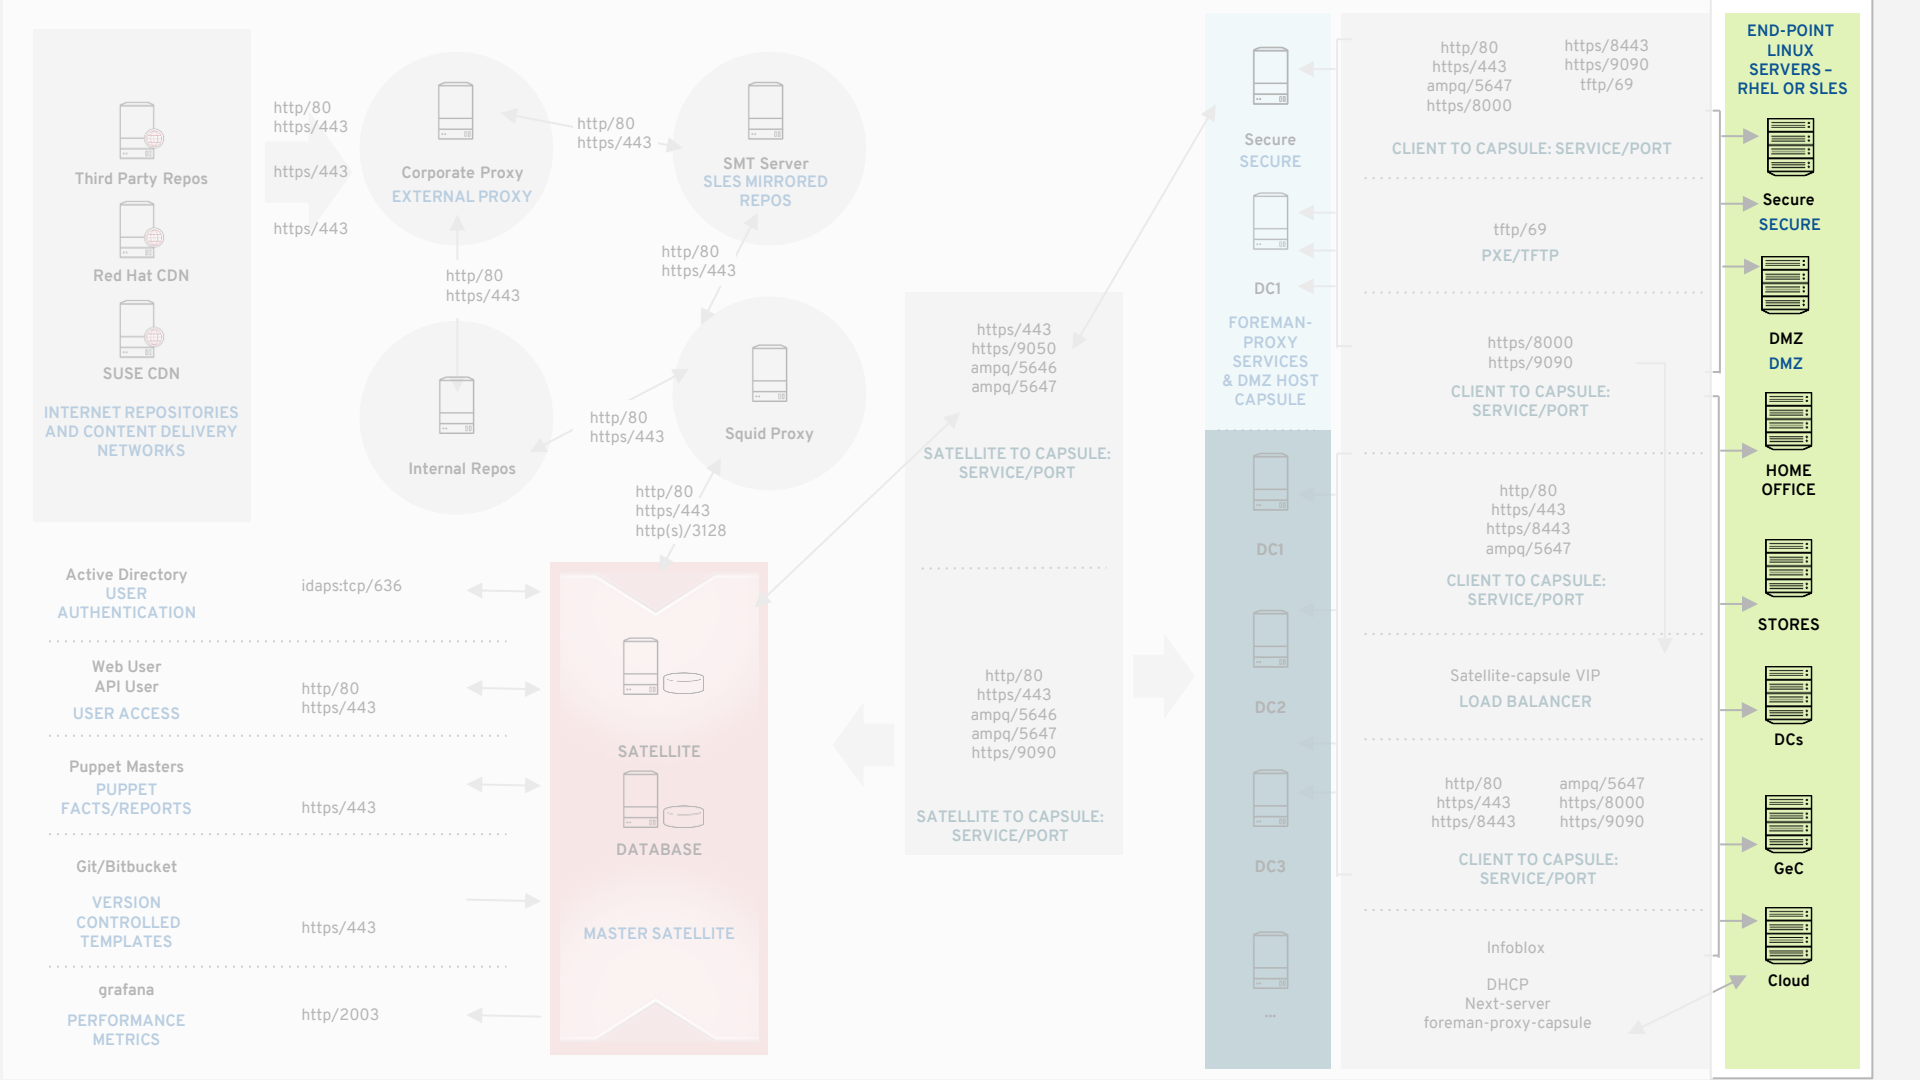
Third Party Repos
Red Hat CDN
SUSE CDN
INTERNET REPOSITORIES AND CONTENT DELIVERY NETWORKS
http/80
https/443
https/443
https/443
Corporate Proxy
EXTERNAL PROXY
http/80
https/443
SMT Server
SLES MIRRORED REPOS
http/80
https/443
http/80
https/443
Internal Repos
http/80
https/443
Squid Proxy
http/80
https/443
http(s)/3128
Active Directory
USER AUTHENTICATION
idaps:tcp/636
Web User
API User
USER ACCESS
http/80
https/443
Puppet Masters
PUPPET FACTS/REPORTS
https/443
Git/Bitbucket
VERSION CONTROLLED TEMPLATES
https/443
grafana
PERFORMANCE METRICS
http/2003
SATELLITE
DATABASE
MASTER SATELLITE
https/443
https/9050
ampq/5646
ampq/5647
SATELLITE TO CAPSULE: SERVICE/PORT
http/80
https/443
ampq/5646
ampq/5647
https/9090
SATELLITE TO CAPSULE: SERVICE/PORT
Secure
SECURE
DC1
FOREMAN-PROXY SERVICES & DMZ HOST CAPSULE
DC1
DC2
DC3
...
http/80
https/443
ampq/5647
https/8000
https/8443
https/9090
tftp/69
CLIENT TO CAPSULE: SERVICE/PORT
tftp/69
PXE/TFTP
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
http/80
https/443
https/8443
ampq/5647
CLIENT TO CAPSULE: SERVICE/PORT
Satellite-capsule VIP
LOAD BALANCER
http/80
https/443
https/8443
ampq/5647
https/8000
https/9090
CLIENT TO CAPSULE: SERVICE/PORT
Infoblox
DHCP
Next-server
foreman-proxy-capsule
END-POINT LINUX SERVERS – RHEL OR SLES
Secure
SECURE
DMZ
DMZ
HOME OFFICE
STORES
DCs
GeC
Cloud

RED HAT®
SATELLITE

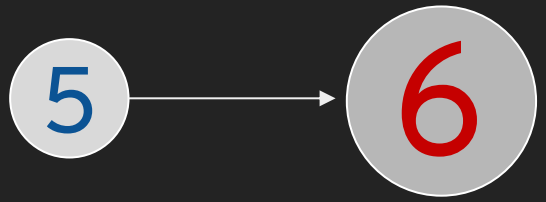

CUSTOMIZATIONS

Hiera customizations
Apache overrides
Limits files
Kernel tunables
Postgres
Logrotate
Pulp concurrency

HIERA CUSTOMIZATIONS

/etc/foreman-installer/custom-hiera.yaml

```
apache::purge_configs: false

apache::mod::passenger::passenger_max_pool_size: 150
apache::mod::passenger::passenger_max_request_queue_size: 100
apache::mod::passenger::passenger_stat_throttle_rate: 120
apache::mod::passenger::passenger_max_requests: 1000
apache::mod::passenger::passenger_min_instances: 6
apache::mod::passenger::passenger_start_timeout: 90
apache::mod::passenger::passenger_max_preloader_idle_time: 0

apache::keepalive: "On"
apache::max_keepalive_requests: 0
apache::keepalive_timeout: 5

apache::mod::prefork::startservers: 8
apache::mod::prefork::minspareservers: 5
apache::mod::prefork::maxspareservers: 20
apache::mod::prefork::serverlimit: 1024
apache::mod::prefork::maxclients: 1024
apache::mod::prefork::maxrequestsperchild: 4000

katello::num_pulp_workers: 16
~
~
-- INSERT --
```

APACHE OVERRIDES

/etc/httpd/conf.d/zzz-custom-overrides.conf

```
<IfModule mod_passenger.c>
  PassengerMaxPoolSize 150
  PassengerMaxRequestQueueSize 100
  PassengerStatThrottleRate 120
  PassengerMaxRequests 1000
</IfModule>

<VirtualHost *:443>
  PassengerMinInstances 6
  PassengerStartTimeout 90
  PassengerMaxPreloaderIdleTime 0
</VirtualHost>

KeepAlive On
MaxKeepAliveRequests 0
KeepAliveTimeout 5

<IfModule mpm_prefork_module>
  StartServers 8
  MinSpareServers 5
  MaxSpareServers 20
  ServerLimit 1024
  MaxClients 1024
  MaxRequestsPerChild 4000
</IfModule>
~
~
-- INSERT --
```

LIMITS FILES

/etc/systemd/system/*

```
/etc/systemd/system/httpd.service.d/limits.conf
[Service]
LimitNOFILE=1000000
LimitCORE=15000000

/etc/systemd/system/qdrouterd.service.d/limits.conf
[Service]
LimitNOFILE=1000000

/etc/systemd/system/qpidd.service.d/limits.conf
[Service]
LimitNOFILE=1000000
~
~
-- INSERT --
```

KERNEL TUNABLES

/etc/sysctl.conf

```
fs.aio-max-nr = 1000000
vm.swappiness = 1
kernel.sem = 250        128000   32      128
vm.dirty_background_ratio = 5
vm.dirty_ratio = 10
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100
fs.suid_dumpable = 1
~
~
~
-- INSERT --
```

POSTGRES

/var/lib/pgsql/data/postgresql.conf

```
max_connections = 1000
shared_buffers = 4GB
checkpoint_segments = 32
checkpoint_completion_target = 0.9
effective_cache_size = 16GB
work_mem = 4MB
log_min_duration_statement = 500
~
~
~
~
~
-- INSERT --
```

LOGROTATE

```
httpd
/var/log/httpd/*log

foreman
/var/log/foreman/*.log
        size=500M

Schedule in cron outside of normal logrotate rules
# Out of band logrotate for high volume logs
0  * * * * /usr/sbin/logrotate /etc/logrotate.d/foreman >/dev/null 2>&1
0 05 * * * /usr/sbin/logrotate /etc/logrotate.d/httpd   >/dev/null 2>&1
~
~
-- INSERT --
```

PULP CONCURRENCY

(Reduced Published Times)

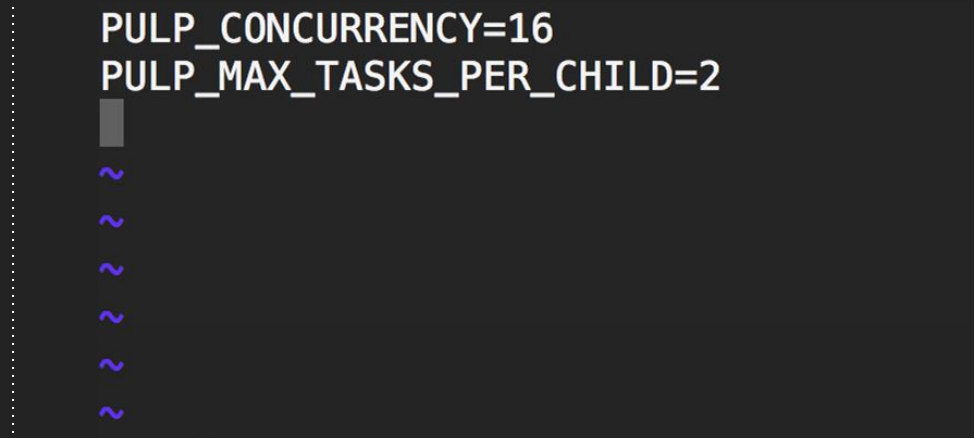

Managing Satellite Clients at Scale

USING PUPPET TO AUTOMATE CLIENT SIDE OPERATIONS

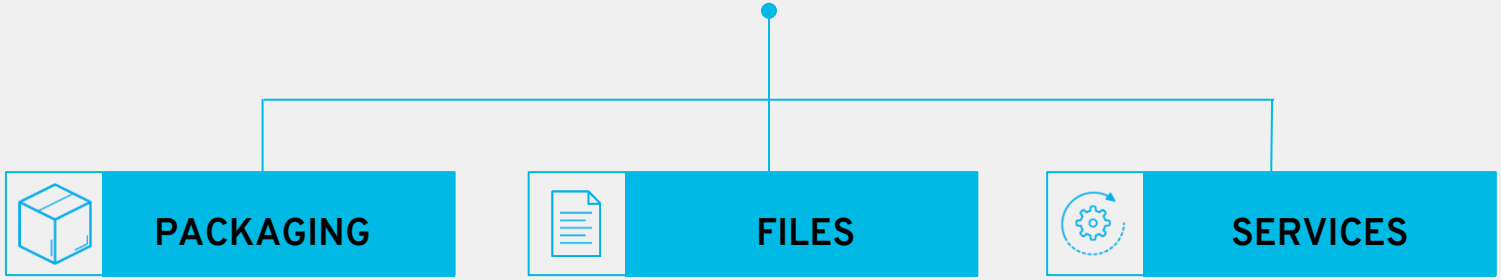

USING PUPPET ENTERPRISE TO ROLLOUT IN SCALE

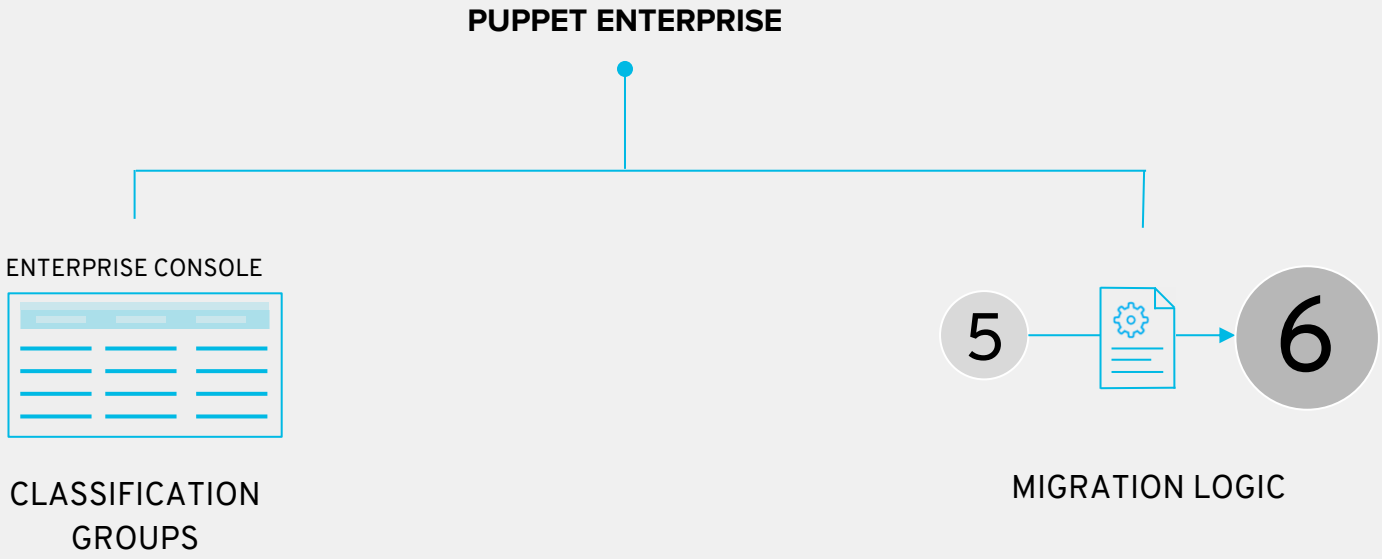

MODIFIED PUPPET MODULE TO RUN MULTIPLE LINUX CLIENTS

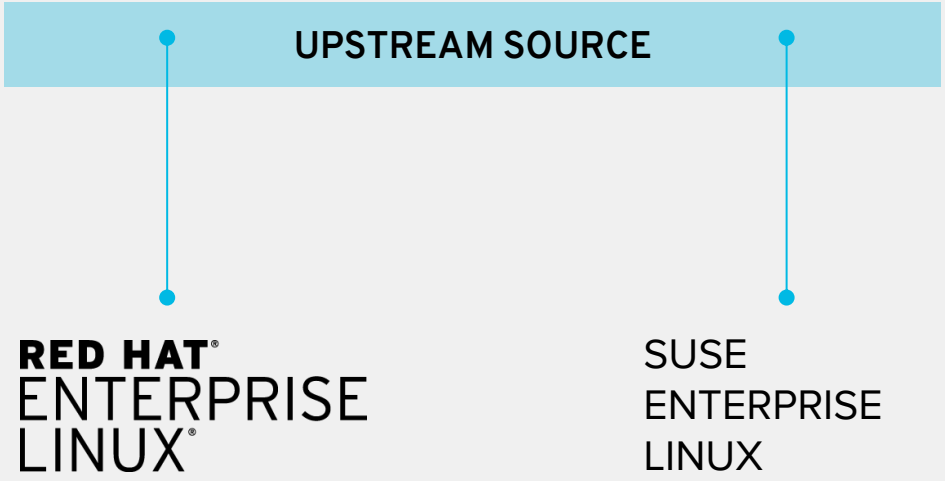

HOW SCALE IMPACTED ROLLOUT

WHERE: **Started in one location**

WHAT: **Touched 60,000 boxes**

WHEN: **Delivered on an aggressive timeline**

HOW: **Rolled out successfully**

Operational Improvements

Walmart had challenges, but they turned them into efficiency gains with customization

BUSINESS CHALLENGE
Need to get clients migrated in a reasonable amount of time. This required automation.

Added automation to roll out to a huge number of clients, multiple locations, small team

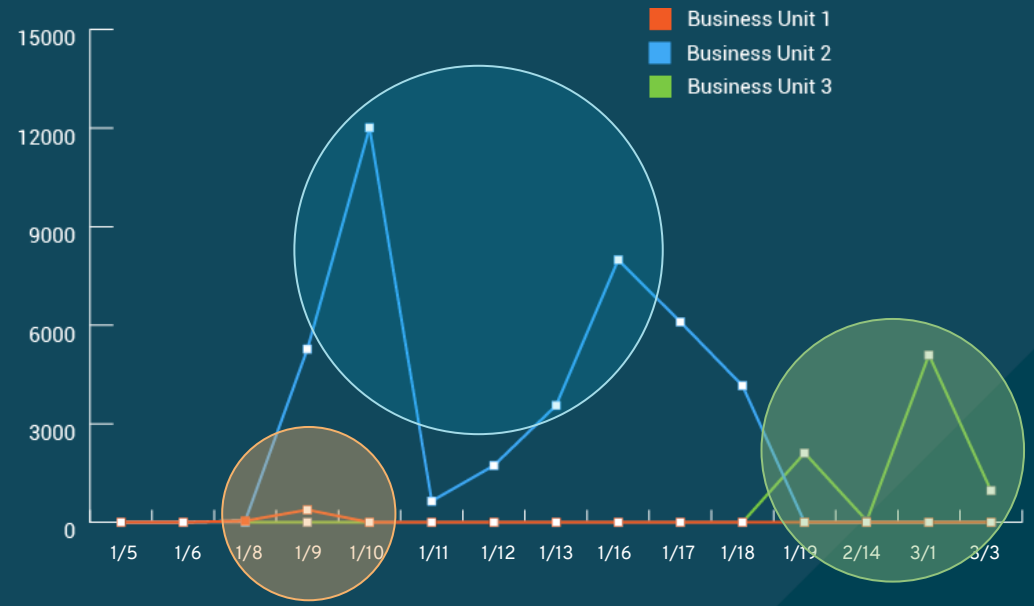

redhat.

BUSINESS CHALLENGE

During client registration, Satellite defaults out of the box are not optimal for large scale, and they get maxed out.

Added randomization to even the distribution

BEFORE

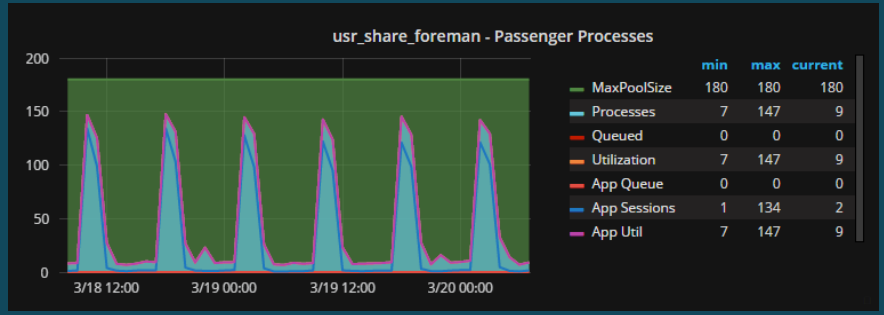

AFTER

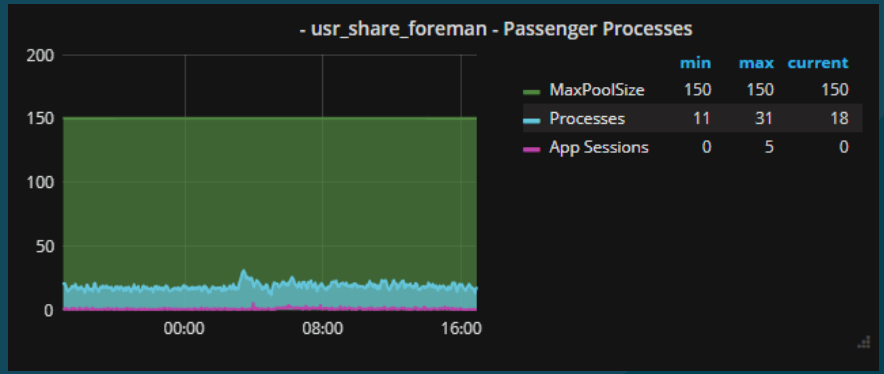

redhat.

BUSINESS CHALLENGE
Users need new content available faster.

Performed customizations to dramatically reduce publish times

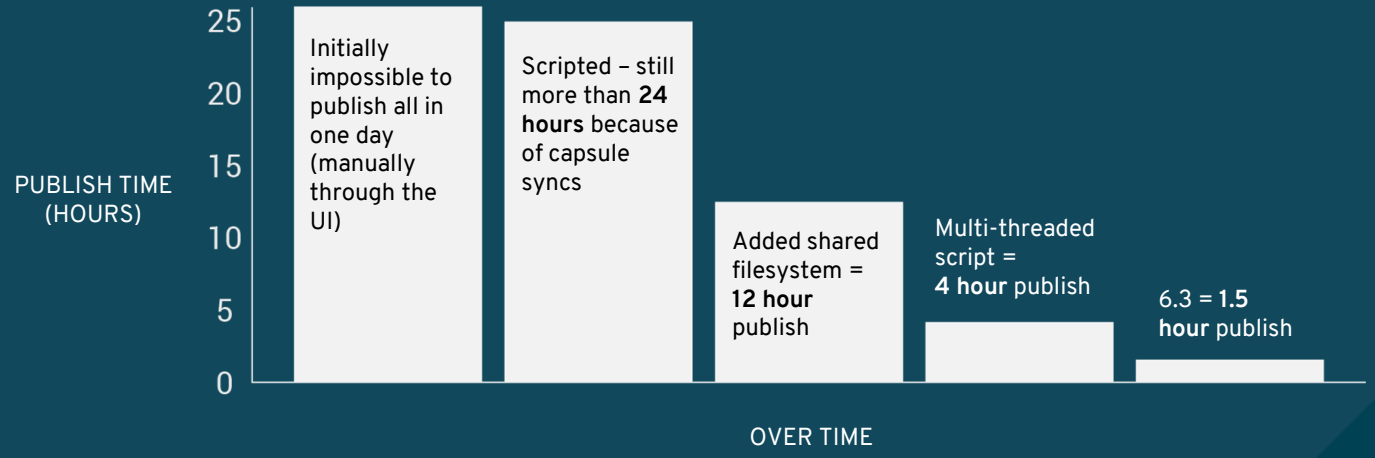

redhat.

BUSINESS CHALLENGE
Walmart needed version control. Template changes from one user changed the template for everyone.

Moved to fully version-controlled template process to improve workflow

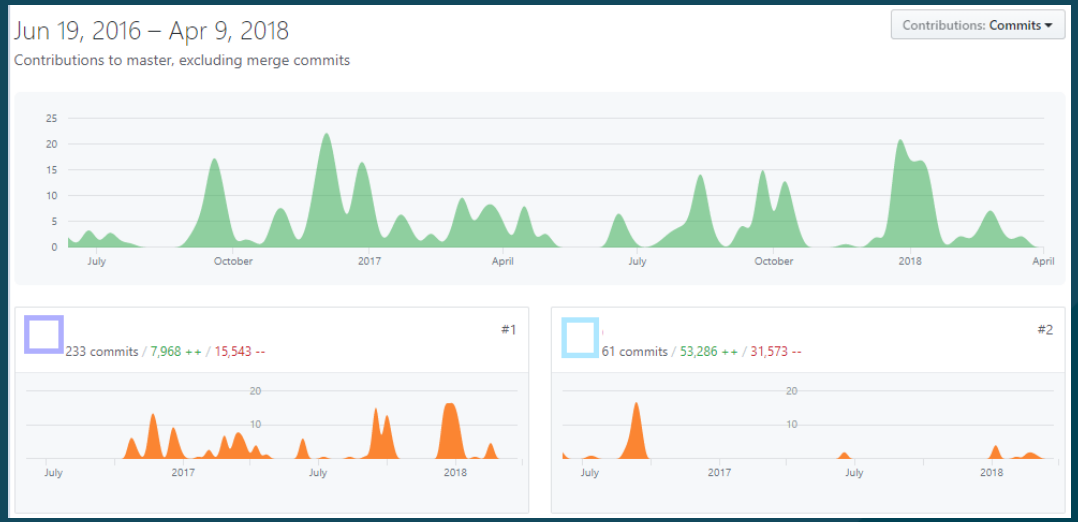

redhat.

BUSINESS CHALLENGE
Need visibility into the operational health of the infrastructure

METRICS

collectd > graphite > grafana
Ansible playbook to redeploy grafana dashboards

IMPORTANT METRICS TO COLLECT

standard cpu/memory/disk/network
foreman tasks active (per task type)
foreman tasks per minute (per task type)
katello event queue
qpid queue depths
dynflow plans
dynflow orphaned execution plans
pulp tasks running
pulp tasks state
pulp repo count
postgres connections
postgres query length
apache processes
passenger processes
apache scoreboards
capsule apache busy servers

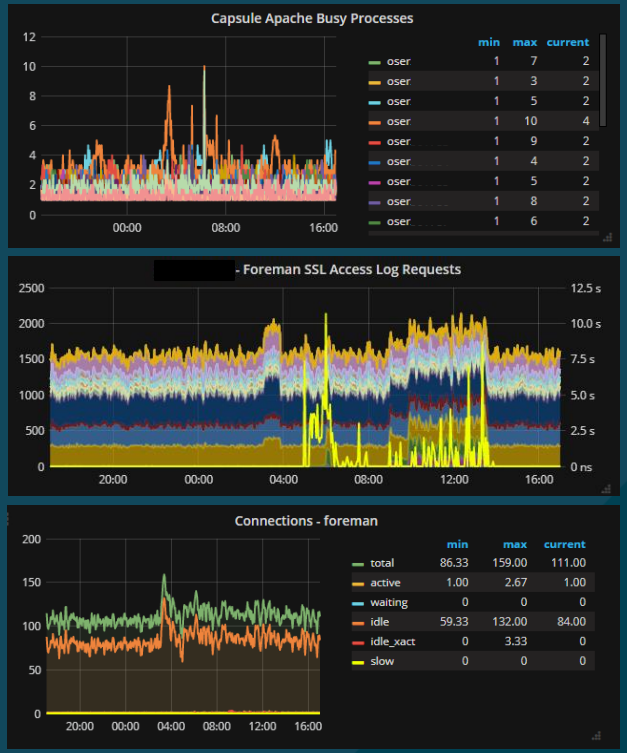

SCALING SUCCESS

Scalability

Build capacity for IT of any size

Collaboration

Share your ideas and optimizations

Infrastructure

Create a reference architecture then customize

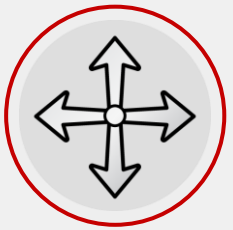

Migration

Know your environment and track key metrics

THANK YOU


plus.google.com/+RedHat


facebook.com/redhatinc


linkedin.com/company/red-hat


twitter.com/RedHatNews


youtube.com/user/RedHatVideos