



# SCALABILITY

- Hodicska Gergely
  - Web Engineering Manager at Ustream
- email: [felho@ustream.tv](mailto:felho@ustream.tv)
- twitter: @felhobacsi

# *DEFINING SCALABILITY*

- It is not:
  - Performance
    - Easier to scale
  - HA
- It is the ability to handle growing amount of work in a capable manner
- Not just technology but people and process

# SCALABILITY TYPES

- Vertical
  - Bigger
  - Typically more expensive
  - Sometimes feasible (DB – SSD)
- Horizontal
  - More
  - Typically we need this

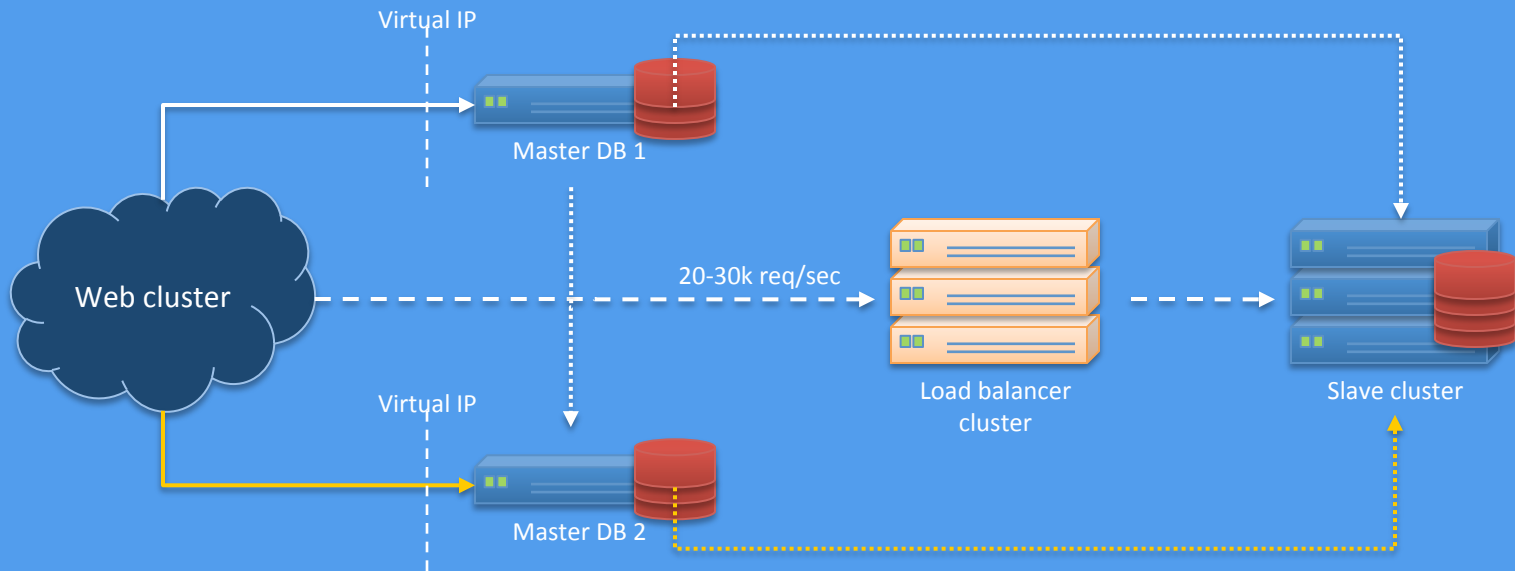
# SCALABILITY RULES

- KISS
- Command and conquer
- Approximate correctness
- Shared nothing

# DATABASE

- Most tough to scale
- Read -> Replication
  - Lag, cascading
- Write -> Sharding
  - App logic, vertical vs. horizontal, key server
- HA: Multi Master Replication
  - DRDB, MMM, MySQL cluster

# U DATABASE – MULTI MASTER MYSQL



—> Normal write

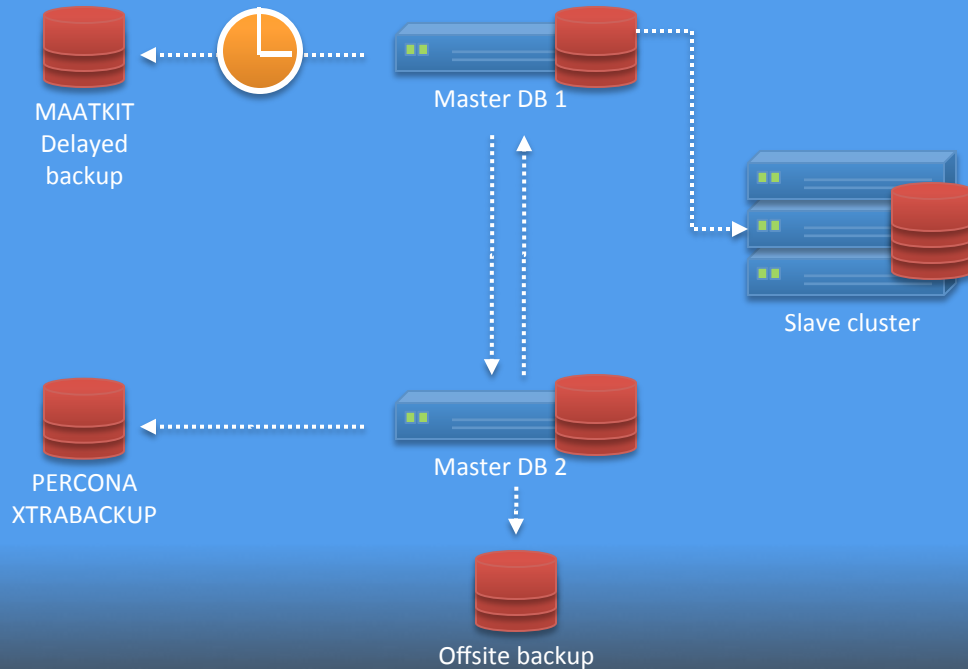
—> Failover write

- - -> Normal read

.....> Normal replication

.....> Failover replication

# DATABASE – BACKUP STRATEGY



- Continuous backup
- Encrypted off site backup
- Delayed replica

# NOSQL

- CAP theorem (availability, consistency, partition tolerance -> eventual consistency)
- Diverse world
- Automatic partitioning, sharding, elasticity
- Transparent for the application
- Extendable without downtime
- Fault tolerant
- Redis, Riak, Voldemort, Cassandra, CouchBase



# CACHING

- Strategies
  - Write-through cache
  - Write-back cache
  - Implicit/explicit invalidation
- Consistent hashing
- Restart

# MEMCACHE

- Local vs. remote
- LRU
- Storing lists
- Versioning
- Race condition (cas)
- Object size

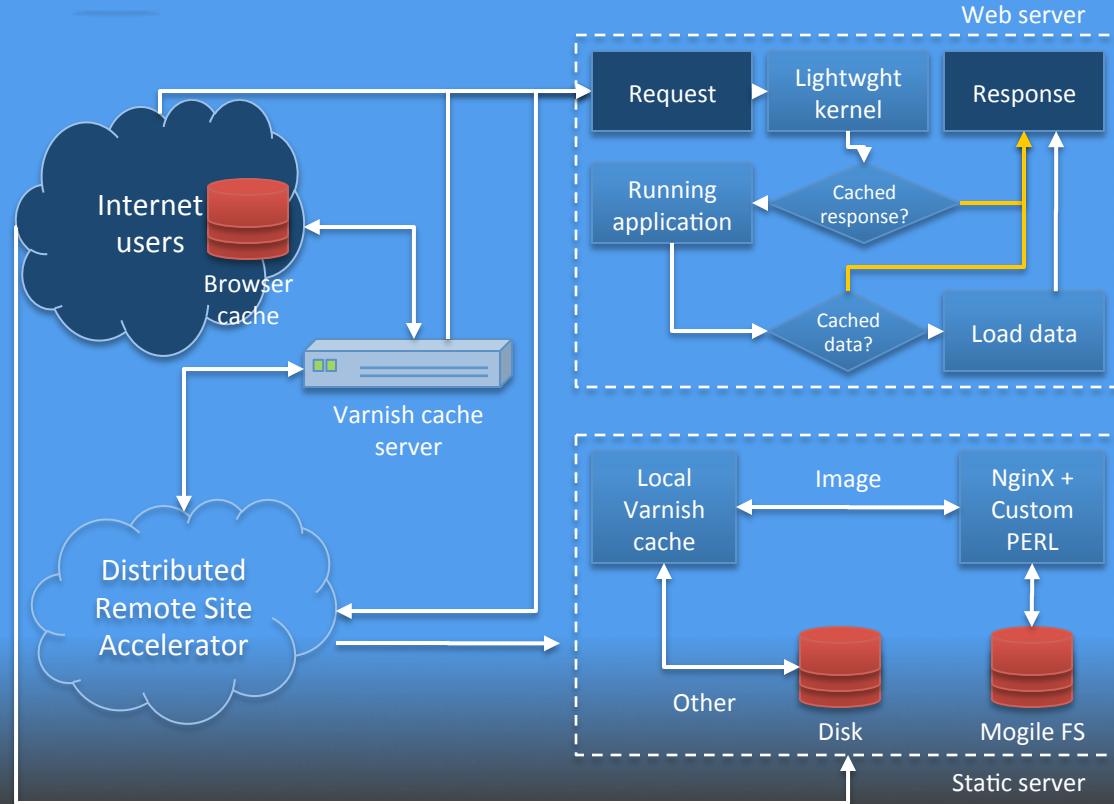
# *APP LEVEL CACHE*

- Code
- Shared memory
- APC, Ehcache
- (Opcode cache)

# HTTP LEVEL CACHING

- Cache-Control header
  - Local vs. proxy
- Static versioning
- Huge expire time

# CACHING



## Highlights

- DRSA: globally distributed reverse proxy to serve the content to the users from a geographically close server
- Cache servers: we try to cache as much requests on these servers as possible to offload our web cluster
- Browser cache: all of our static assets are automatically versioned for optimal serving (using huge expire times)
- Application level: caching whole pages to avoid running all the code or the pieces of data to offload the database
- Our framework automatically package and compress the JS and CSS files to reduce the number of HTTP requests



# STATIC CONTENT

- NFS (mount problems)
- DB
- MogileFS
- Authentication, access control
  - Perlbal
- FS limitations
- Low hit ratio

# *LOAD BALANCING*

- Dedicated hardware vs. software based
  - Price
- HA proxy
- Nginx (HTTPS termination)
- LVS (direct routing)
- DNS loadbalance, Anycast (multi site)
- Layer 4 vs. 7

# SESSION

- Sticky (shared)
- Centralized
  - DB / NoSQL
  - Memcached
- Cookie
  - Sysop will like you



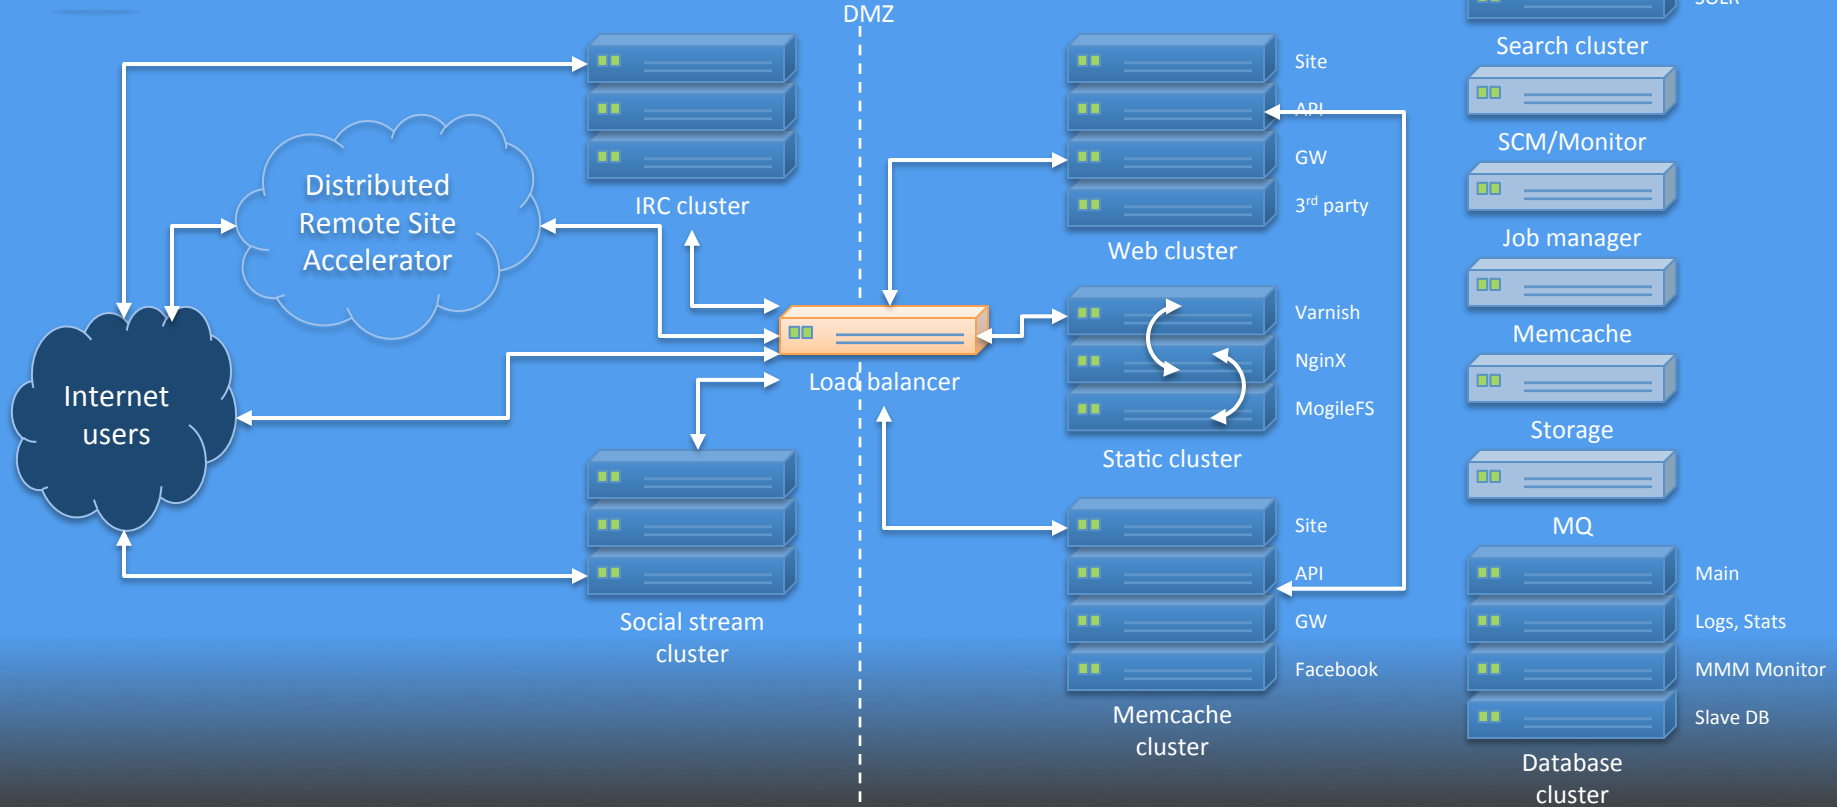
# ASYNCHRONOUS OPERATIONS

- Decoupling
- Capacity handling
- Node.js
- Jobs
  - Gearman
- Message queue
  - Q4M, RabbitMQ, ZeroMQ

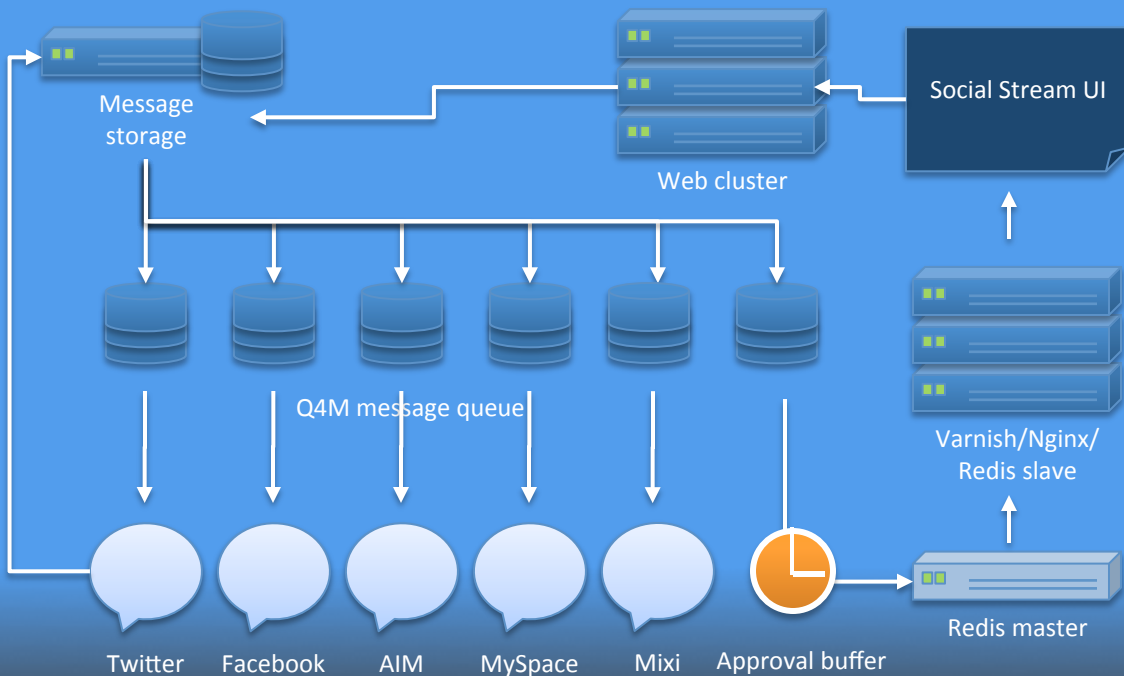
# CHALLENGES OF THE WEB STACK

- Web requests:
  - 200 requests / sec / server (peak)
- Cache server requests:
  - 10,000 requests / sec / server (peak)
- Social stream requests:
  - 15,000 requests / sec / server (peak)
- Database requests:
  - 25,000 requests / sec / server (peak)

# WEB ARCHITECTURE



# SOCIAL STREAM: UNDER THE HOOD



## Highlights

- Generated 2.5M visits in the last 30 days (the 20% of this is new visitor)
- 0.8-1.1M messages per day
- Justin Bieber has ~230k messages per day
- Jonas Brothers concert: 10k messages per minute in peak
- Peak:
  - 5k new connection / sec
  - 15k requests / sec
  - 600 Mbit / sec

# *DEVELOPMENT BEST PRACTICES*

- Continuous integration
  - Automated builds
  - Unit tests
  - Acceptance test
- TDD, (ADD, FDD ;))
- Abstract branching (feature switch)
- Code review, pair programming, topic experts
- DevOps culture

# DEVOPS TOOLING

- Provisioning
- Configuration management (cfengine, chef, puppet)
- Application deployment (capistrano, fabric)
- Orchestrator (mcollective)
- Monitoring (system/application level)
  - Nagios, Munin, Cacti, Graphite etc.
- Supervisors (monit, god)
- Log management/analysis

# DEVELOPMENT BEST PRACTICES

- Visualizing
  - Graphite (StatsD, logster)
  - Custom dashboards with KPIs, alerting
- Runtime vs. build time
- Automated code deployment
- Load testing (automated better)
- MVC

# CONFIGURATION MANAGEMENT

- Automation
- Versioning
- Accountability
- Chef, Puppet
  - Same local dev environment



# *CLOUD ENVIRONMENTS*

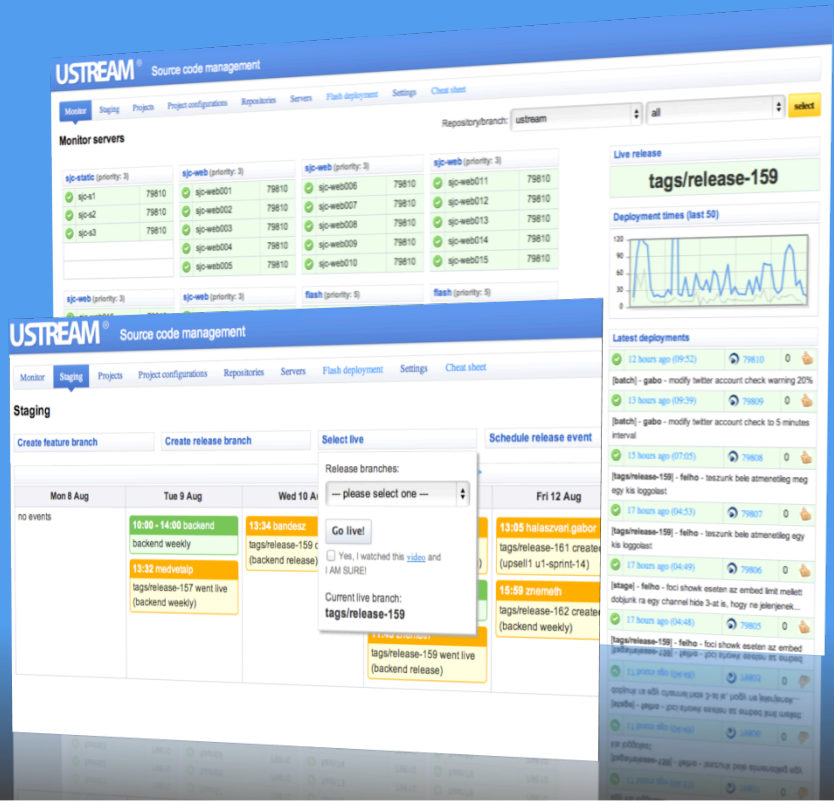
- Amazon, Google App Engine etc.
- Early stage
- Cost savings
- Backup for peaks

# ORGANIZATIONAL BEST PRACTICES

- Architect board
- Scrum of scrum
- Product board
- WESK (who else should know)
- Internal demos

# SOURCE CODE MANAGEMENT

- Using git for most of the projects
- Code is deployed to more than 120 servers (parallel deployment)
- Custom SCM tool which automates the process
- Different environments: development, staging, releases, live
- Devops culture: developers, sysops and QA work closely together
- Release and rollback policy



The screenshot displays the Ustream Source Code Management interface. At the top, there's a navigation bar with tabs for Monitor, Staging, Projects, Project configurations, Repositories, Servers, Flash deployment, Settings, and Cheat sheet. Below this, the 'Monitor servers' section shows a grid of server status for various environments like 'sjo-static', 'sjo-web', and 'flash'. The 'Staging' section features a calendar view of deployments and a 'Go live!' dialog box for 'tags/release-159'. The right side features a 'Live release' section for 'tags/release-159' with a deployment time graph and a list of 'Latest deployments' with their respective commit messages and times.



# MONITORING

## Highlights

- Proprietary dashboard to oversee key system performance charts in one location
- Real-time information about streaming related servers, web/cache servers, database servers
- Summary of the Nagios checks
- Ability to roll back for historic charts
- Provides shortcut to system tools

## Tools

- Munin: Several custom plugins
- Cacti: Mainly for network devices
- Nagios: More than 1200 checks, Active checks
- Monit: Ensuring that a given process runs and it doesn't consume too much resources, Active checks
- Query watchdog: Automatically stops and reports excessive read queries





# WHERE TO IMPROVE

- English
- Enjoy programming
- Soft skills (communication, team working, presentation, cooperation, management, leadership, time management etc.)
- Agile development methods (Scrum, Kanban)
- Continuous learning (blogs, books, conferences, code)
- Craftsmanship
  - Clean Coder, Agile Software Development, Martin Fowler books / signature series, The Mythical Man-Month, Code Complete, The Pragmatic Programmer, Peopleware, The Passionate Programmer



# WHERE TO IMPROVE

- Network programming, protocols
- Algorithms
- DHT
- Database
- Big Data (Hadoop, HBase, Hive/Pig, BI tools etc.)
- API design (REST, SOAP, oAuth etc.)
- Playing with open source tools of big companies (Twitter, FB, LinkedIn)
- Blogging, taking part in open source projects
- Learning different type of programming (e.g. functional)




*THANK YOU*

Questions?



*NOW PLAYING*  
SCALABILITY – BME

 32

 May 7, 2012

**LIVE**