

CASE STUDY

Traffic is coming ! “OMG” moments

*How Distinct helped
evoMAG.ro handle Blackfriday
2013 traffic spike*

Summary:

1. Who
2. The Challenge
3. The Work
4. Aftermath

1. Who :: Bogdan Belu - Distinct

Bogdan Belu - www.bogdanbelu.com

- 19+ years of technical and business experience (network administration, programming, datacenter services)
- since 2001, CEO of Distinct (www.distinct.ro), managed services and datacenter provider under "plus through responsibility" motto
- helps clients handle their mission critical Internet infrastructure

www.distinct.ro

Clients



1. Who :: evoMAG



online retailer with IT-related products focus

2013 total revenue: 15.5 mil EUR

BIG Blackfriday promotional sale on Nov, 22, 2013 (up to 99% price cuts)

2. The Challenge

“keep the www.evomag.ro site 100% up during the Blackfriday 2013 event”

2. The Challenge :: Context

- ❑ a warehouse full of products waiting to be sold
- ❑ uptime problems in the previous years
- ❑ same web application (more or less) as the previous years
- ❑ better marketing campaign from evoMAG team compared to the past
- ❑ marketing and inventory disaster if something goes wrong





2. The Challenge

Distinct's role in this project:

- datacenter services (servers, networking)
- managed services (system administrators)

Gladly accepted the challenge since is consistent with our mission:

“Distinct provides infrastructure and maintenance of Internet services and applications. Distinct's offer include datacenter hosting, managed services and SaaS. The success of our clients depends on the proper functioning of computer applications such as content services, e-commerce, internet banking or email.”

3. The Work :: Analysis

- new client for Distinct
- started work about 40 days before the event (late start)
- got historic metrics from the previous Blackfridays
- got the traffic estimates from evoMAG's marketing departament (number of visitors, number of orders, estimated user path in site, etc)
- analyzed the site thoroughly (technical point of view)
- efficient budget because of the low margins (more of a marketing event)

3. The Work :: Targets

- translated marketing estimates into technical metrics (HTTP hits per site feature, bandwidth, etc)
- better to overestimate than sorry
- further translated technical metrics into an infrastructure proposal (network, servers, software stack)

Orders: prepared for 50 orders/s (peak)

Example for the **Homepage**:

- estimated the number of visitors per day: **600 000**
- estimate the peak number of unique visitors per hour: **200 000**
- total size of homepage: **1 MB**
- total needed bandwidth for homepage for one user: **8 Mbps**
- HTTP hits per homepage for one user: 100 (first visit, no browser caching)
- average unique visitors per second (peak): **55.56**
- estimated hits per second: **5555.56**
- estimated bandwidth needed: **444 Mbps**
- after that, overestimated about **2.5x** in order to handle instant peaks/stay safe

3. The Work :: Roles

- **Distinct** (datacenter and managed services): systems and network
- **Customsoft** (the web application developer): special Blackfriday web pages, code optimizations, application insights
- **evoMAG** (the client): marketing data, constant feedback and rapid approvals

team work and
communication
is essential
(really !)

neither two of
the parties
involved could
handle this
without the
third

approach:
break the 'big
problem' into
small pieces
and solve
them

3. The Work :: Technical Keypoints

- URL analysis (load impact, predicted usage, usefulness) and classification, separate static (eg: product description, content page) and dynamic content (eg: shopping cart) into 'content classes'
- explicit URLs in order to ease classification
- **took the responsibility for URLs** (hardcoded caching and routing rules in load balancer), more than 1000 lines of configurations
- suggested to temporary eliminate unnecessary stuff for Blackfriday or make it client side, if possible (popups, cookies, background processes - crons not critical)
- request small but with big load impact software modifications whenever possible (due to short time frame)
- tried not to modify the existing setup core too much (to avoid the 'joys' of refactoring in such a short time), **built around it**
- move most of the dynamic bits on **AJAX**

3. The Work :: Technical Keypoints

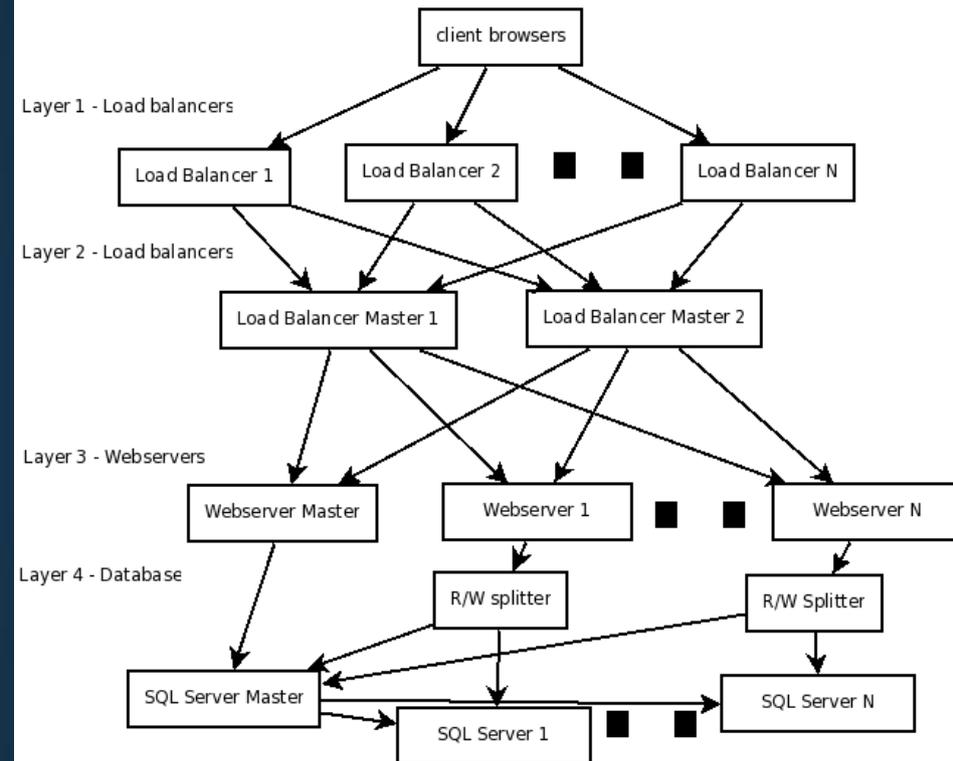
Layer 1: Caching/Load Balancing

- those are the the systems client browsers connect to
- 16 IPs in DNS for www.evomag.ro at peak (load balancing via DNS)
- roles: caching, hits spikes load bearing
- scalable (easily add more on demand)
- HTTP caching, lots of (> 90% hit rate)
- automated health checks of the next layer

Layer 2: Caching/LoadBalancing/Routing

- 2 servers (active/active configuration)
- cache aggregation
- permitted quick/realtime request routing changes (automatic or manual)
- automatic health checks for the next layer (webservers) and exclude the 'bad' servers from production for a while (ex: 30s)
- even if the webserver behind is unavailable deliver cached content (when appropriate, we frozen static pages modifications for the day)

TOPOLOGY:



3. The Work :: Technical Keypoints

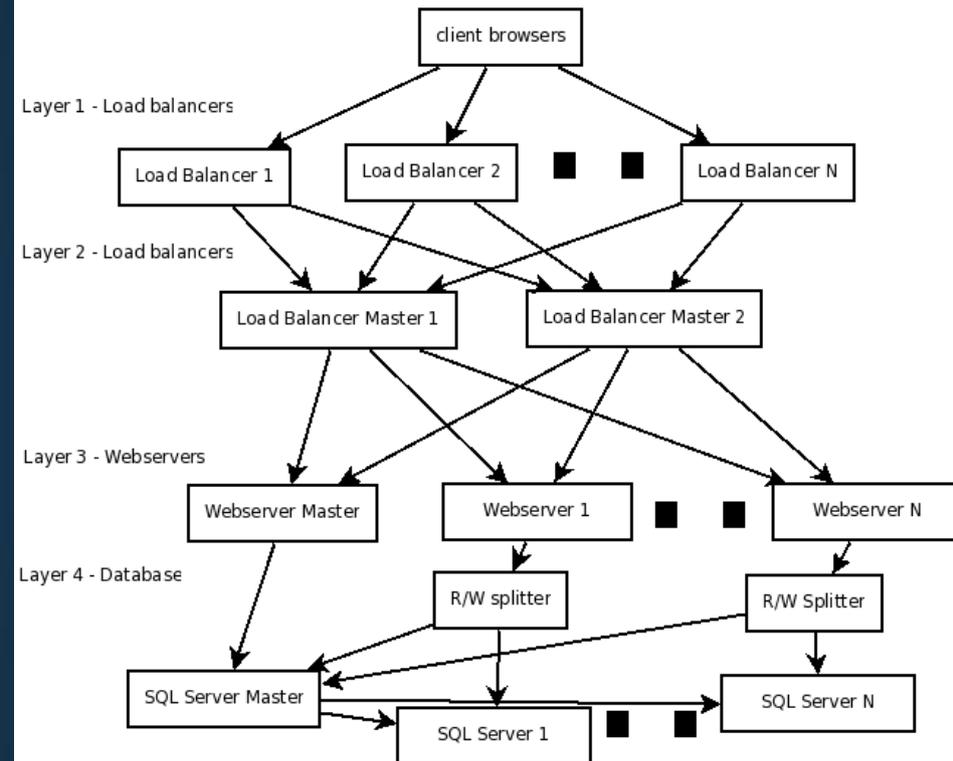
Layer 3: Webservers

- cloned the existing webserver into multiple copies
- route the important requests (eg: orders, stock availability) into the designated master server
- designed a master server for 'sensitive' operations in order to workaround possible stock synchronization problems (because 'out of stock' condition could be triggered in matter of seconds for certain 'hot' items, each item had limited stock)

Layer 4: Database Servers

- used replication + R/W splitting component
- managed to avoid MySQL clustering (risky)
- optimized sql tables usage, temporary eliminate less important ones
- each webserver paired with a SQL server

Topology (continued):



3. The Work :: Technical Keypoints

Hardware/Network:

- used **21 separate servers** (not counting multiple virtual machines on each server)
- **hybrid cloud** approach (evoMAG owned original servers + Distinct dedicated servers + cloud servers available to scale the frontend caching)
- Software Stack: Linux (OS), Openvz (virtualization), Nginx (webserver), Mysql (database), PHP (programming language), memcached (sessions), Varnish (caching/reverse proxy)
- extensive use of **SSDs** for databases
- avoided use of a distributed file system (simpler stack, application permitted that), used periodic sync to clones + automatic request rerouting in case of errors (via load balancers)
- did **load testing** scenarios (we could have used more)

3. The Work :: Technical Keypoints

1. Redundancy:

- DNS (automatic)
- Reverse Proxy/Load Balancers (automatic but first layer semiautomatic)
- Webservers (automatic)
- SQL Servers (automatic, except master SQL server; plan: manually promote replica if failure happens)
- reduce to minimum the SPFs (Single Points of Failure)

2. Scalability:

- on every layer we could add more resources (horizontally)
- except master SQL server (which handled a limited set of requests by design)

3. Security:

- prevent possible security infiltrations masked by the high traffic

4. Monitoring:

- extensive monitoring, 94 key sensors (including error rate checks and automated simulated order of products)
- sms/phone/messenger alerting

3. The Work :: Black Friday Timeline

- started adding servers (traffic build-up, as expected)
- everybody on duty (Distinct, programmers, evoMAG) and in constant communication

Traffic is coming ! “OMG” moments

- last minute glitch in the customized homepage (heavy ajax), masked by using a manual LB rule in a matter dozens of seconds
- exceeded 10 000 users on site about 20 minutes after T0 (source: Google Analytics, 'Right Now' visitors on site)
- about 20 000 HTTP hits per second peak (within estimates)
- servers load ~ 25% (good) with some spikes

T0-4 days

T0-7h

T0 - 00:20, Black Friday (22 Nov 2013)

- launched teaser homepage + newsletter signups
- so many newsletter signups we thought something is wrong (> 100 000)
- tried to hold any last minute modifications of the system (not so successfully)

3. The Work :: Black Friday

În acest moment,

9460

(de) vizitatori activi pe site

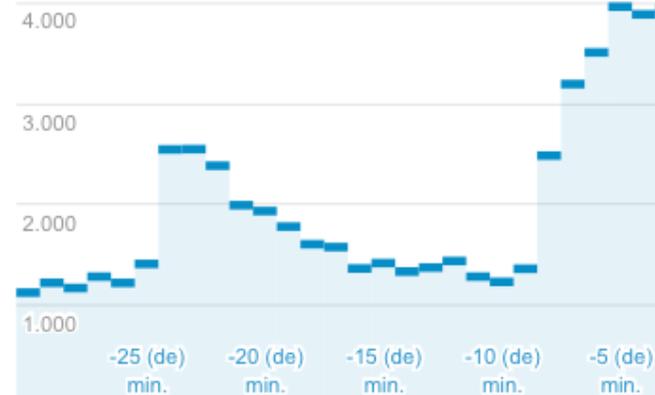
 RECURENT  NOU

74%

26%

Afișări de pagină

Pe minut



Pe secundă

150

100

50

0

0

0

0

0

0

4. Aftermath :: Numbers

T0 + 26 minutes: orders 'golden minute':

- 365 orders / minute
- 10 000 unique visitors on site according to Google Analytics

- T0 + 9h: 65% percent of inventory sold
- T0 + 9h: 145.317 visits
- T0 + 9h: ~ 9000 orders

- T0 + 12h: 1.1 Million EUR sales

Blackfriday totals:

415 000 total visits

1.6 Million EUR sales
(~ 10% of year's revenue)

number visitors growth
compared to the previous
Blackfriday: +66%

newsletter total
subscribers: +35.5% up (as
many as gathered in the
previous 4.5 years)

(data source: evoMAG)

4. Aftermath :: Feedback

- positive feedback from the customer, Mihai Patrascu, CEO, evoMAG:

“I am sending to everybody in the team a sincere THANK YOU !”,

- positive feedback from Internet users (blog comments/forum posts), random sample:



Translation: “Evomag was decent until now, haven't crashed at all. We'll see eMAG by morning”

Source: <http://www.buhnici.ro/2013/11/22/lista-evomag-black-friday-2013/>

4. Aftermath :: Feedback

- Positive mass media feedback, example PRO TV / “I like IT” (special edition about Blackfriday), starting from 4m:40s:



<http://stirileprotv.ro/stiri/ilikeit/secretele-vanzarilor-black-friday-vor-fi-dezvaluite-live-greii-emag-evomag-sambata-la-ilikeit.html>

Q & A

Thanks !

Bogdan Belu, CEO Distinct
Email: bogdan.belu@distinctgroup.net

www.distinct.ro

Thanks also to our
partners



*This presentation (slides + video)
available on www.bogdanbelu.com*

