

TUNING LINUX FOR BIG FIREBIRD DATABASE: 693GB AND 1000+ USERS

Alexey Kovyazin

www.ib-aid.com

Firebird Conference 2019

Berlin, 17-19 October



YOUR PREMIER SOURCE OF FIREBIRD SUPPORT

IBSurgeon



**MOSCOW
EXCHANGE**



Fast Reports
Reporting must be fast!

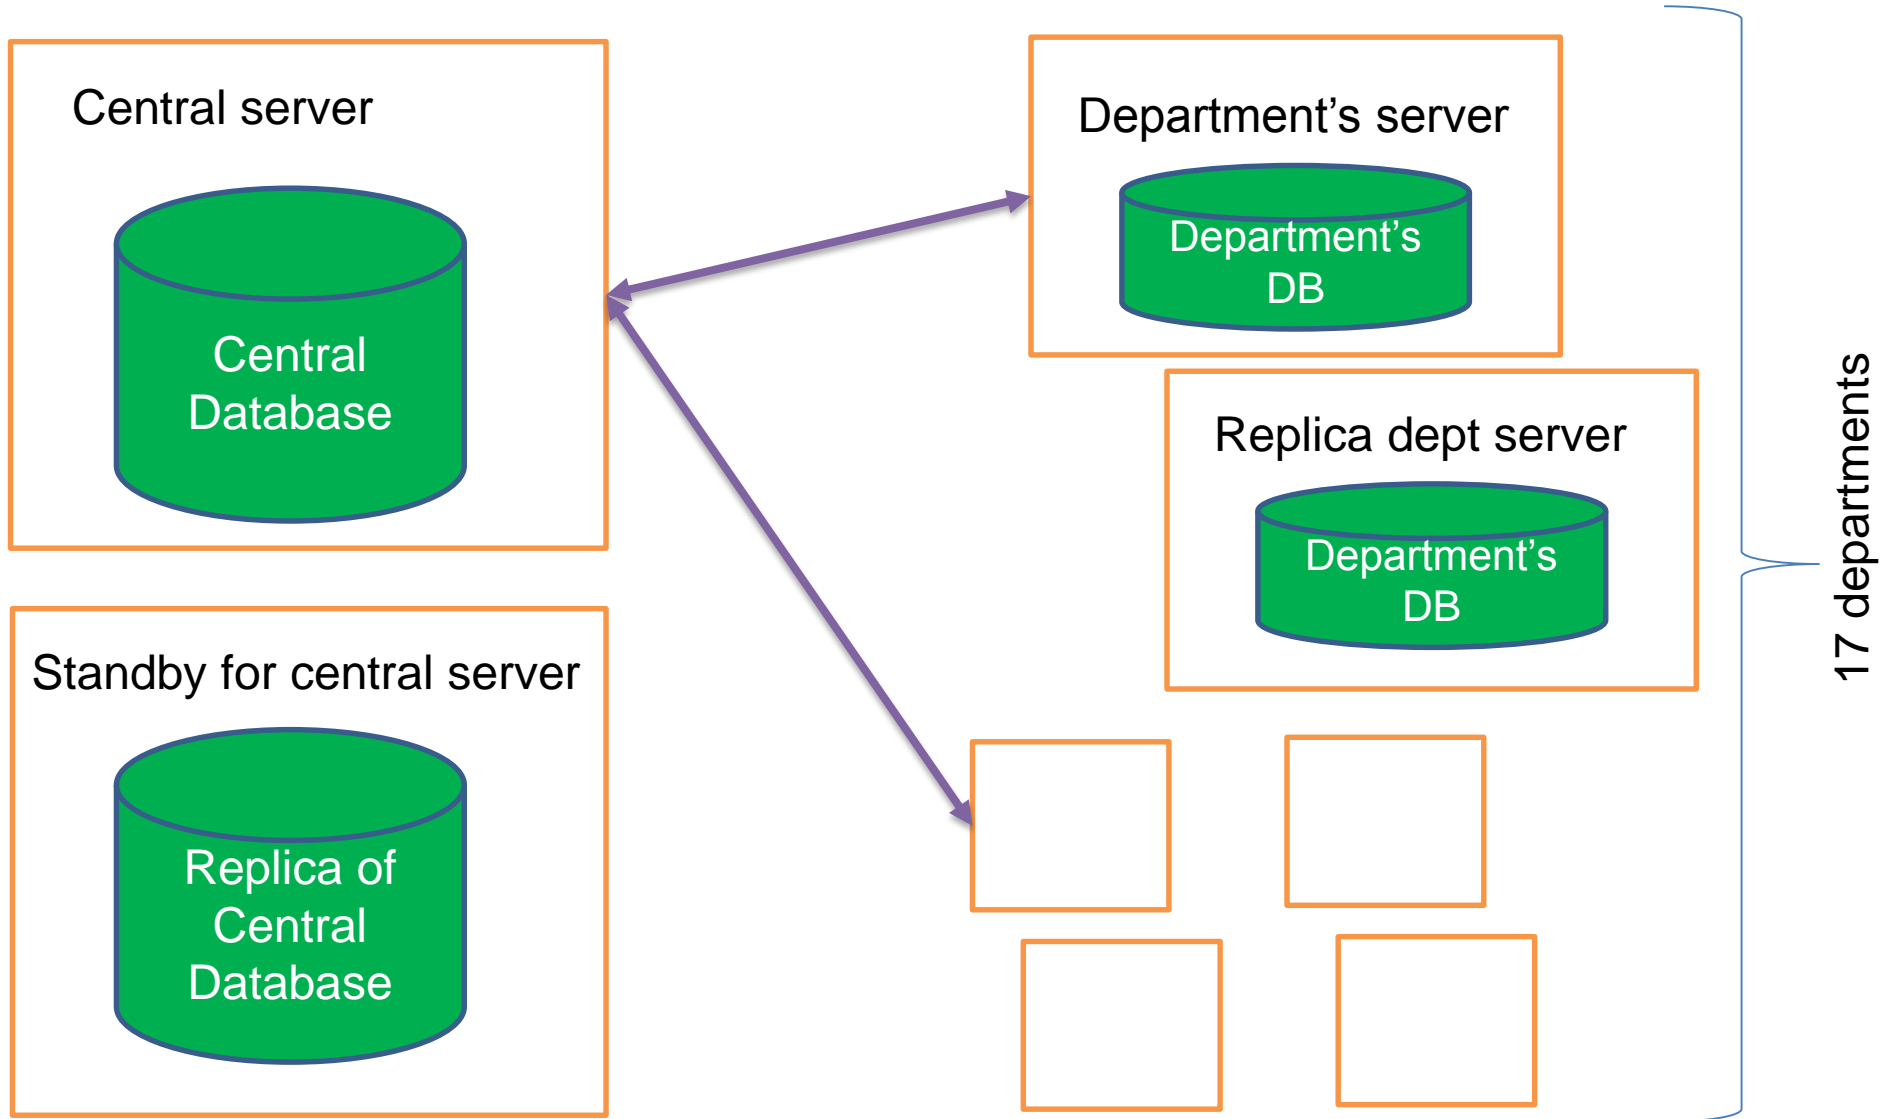


Customer: <http://klinikabudzdorov.ru>



- BudZdorov
- Medical centers and hospitals in Moscow, Saint-Petersburg and major cities in Russia
- 17 departments
- 365 days per year, from 8-00 to 21-00

ERP with Firebird in BudZdrorov



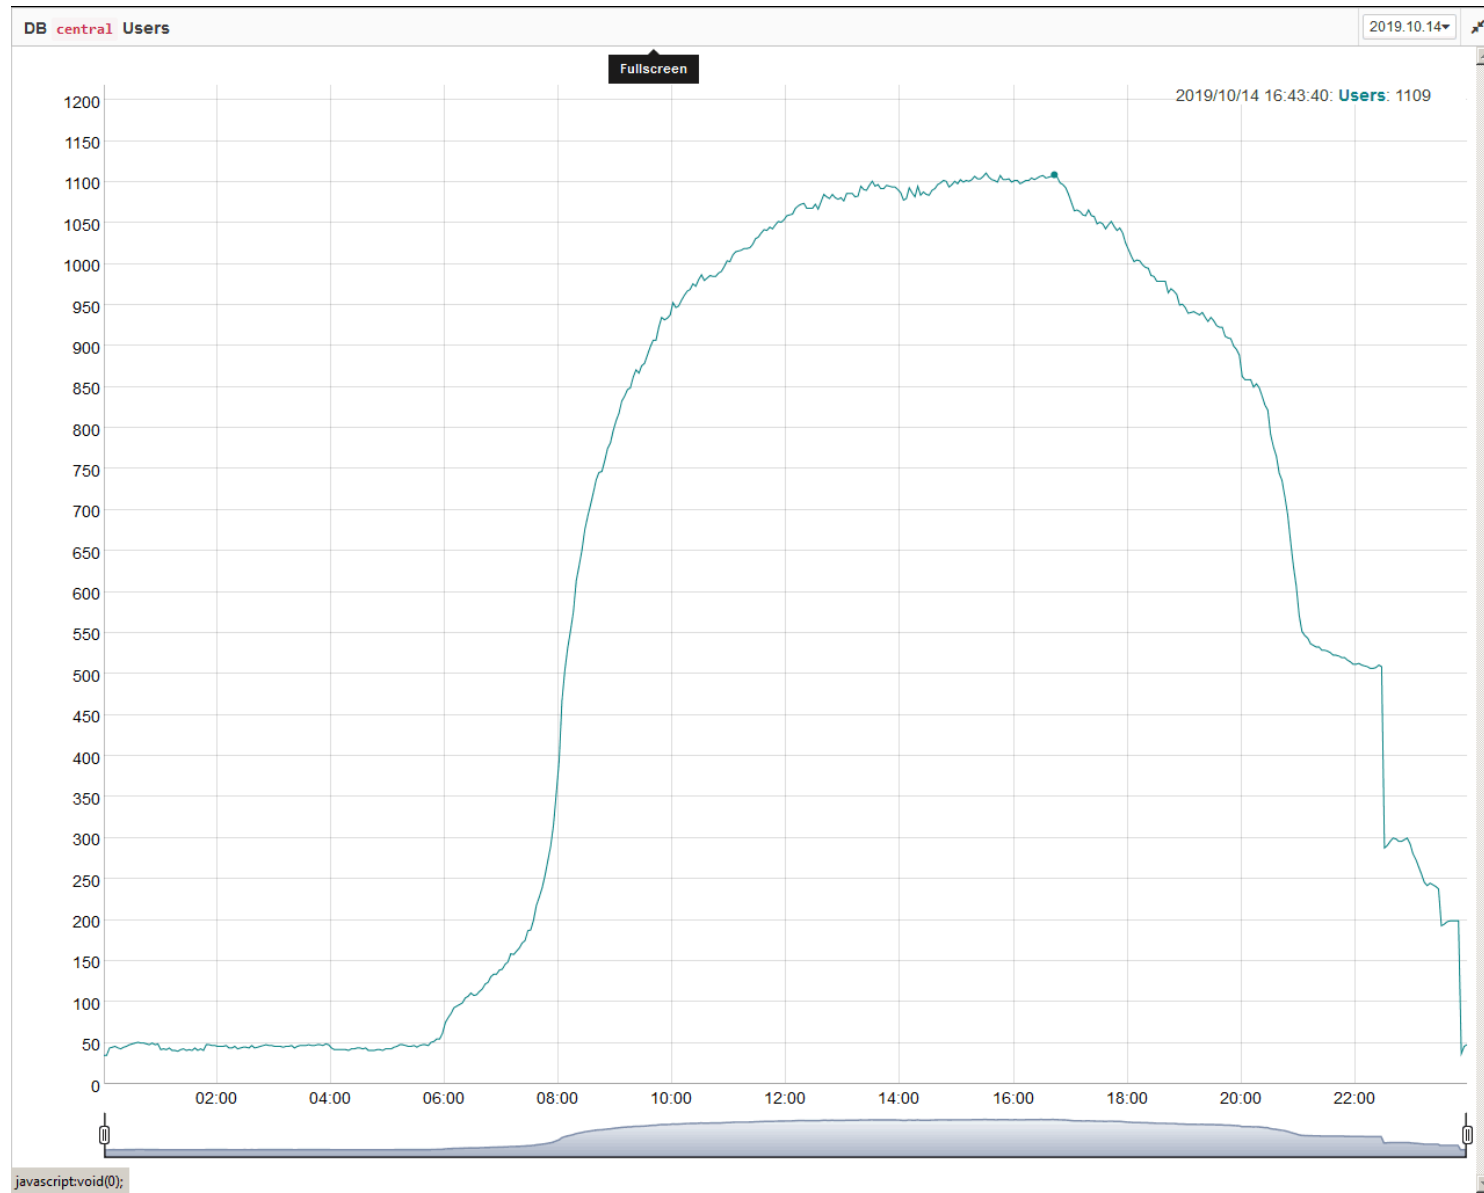
Central database in 2017

- Size = 453 Gb
- Daily users = from 700 to 1800 (peak)
- Hardware server
- OS = Linux CentOS 6.7
- Firebird 2.5 Classic + HQbird
- Client-server, connected through optic with departments
- With async replica on the separate server


...and in 2019

- Size = 665 Gb
- Daily users = from 900-1100 (normal day) to 1800 (peak)
- Hardware server
- OS = Linux CentOS 7.6.1810
- Firebird 3.0 HQbird

Users: normal day 1000-1100



Database growth: 691Gb, +0,5Gb/day



Hardware

- HP ProLiant DL380p Gen8
- 320Gb RAM
- Gen8 2x Xeon(R) CPU E5v2 @ 2.60GHz
 - 24 logical cores with HT
- 4 network card

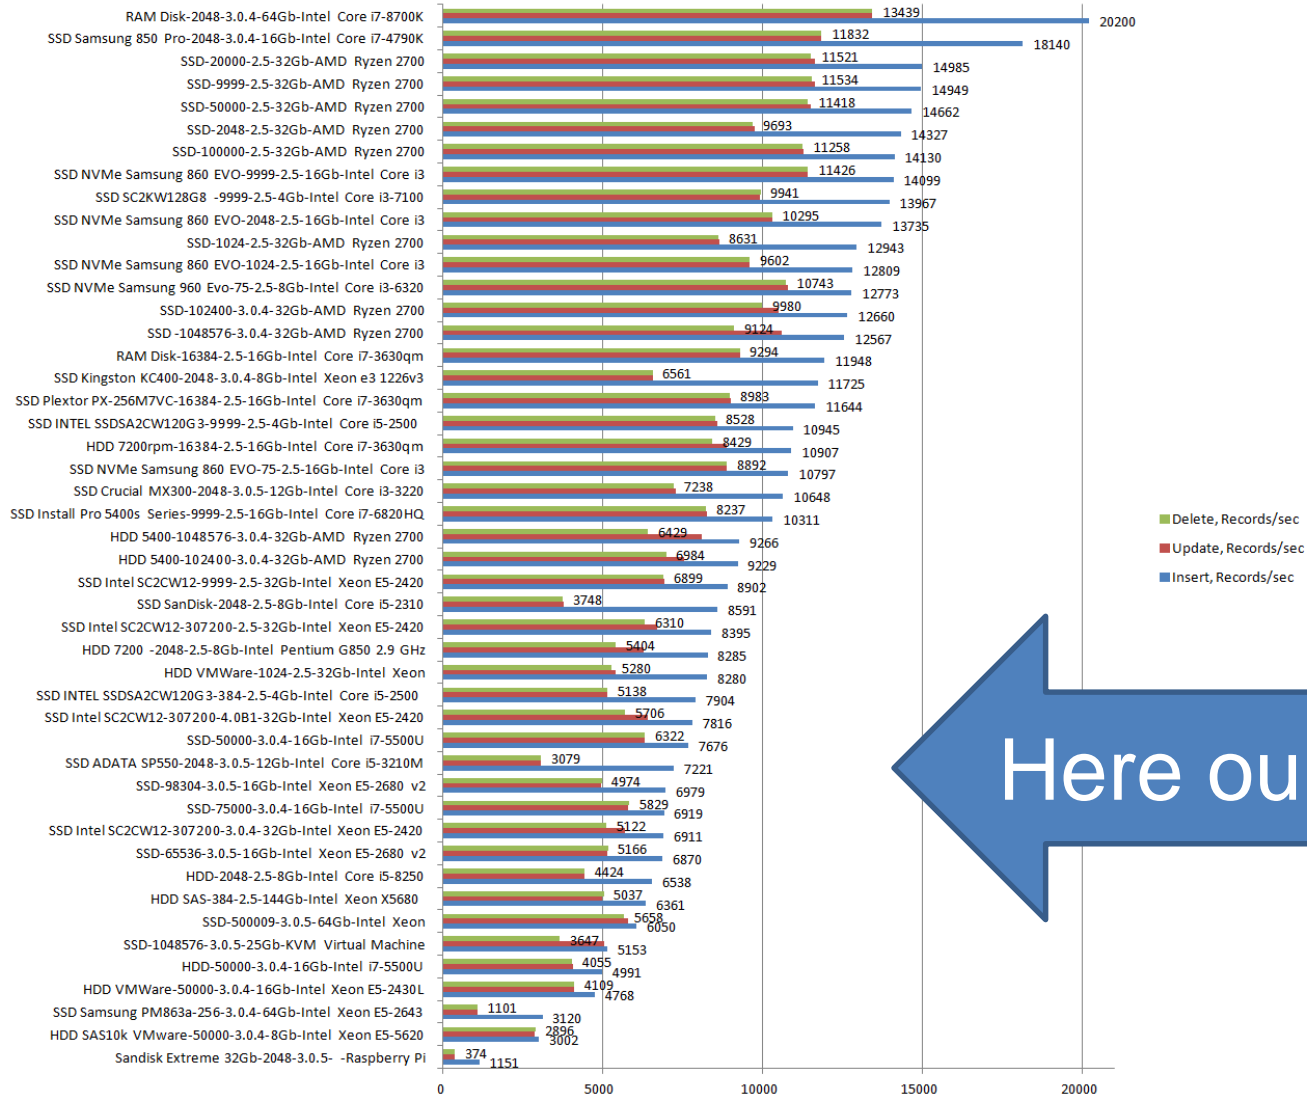
Disks

- sda (745Gb),
- 2 fibre-channel connection to SAN (FCoE):
 - sdb 1.8Tb
 - sdc 1.8Tb

Disk test

- <https://ib-aid.com/dbtest>
- SQL script to INSERT/UPDATE/DELETE of 1 mln records in 1 thread

Hardware is not the best



Here our server

Challenges

- Database grows, hardware not
- Performance degradation during the long periods of uptime
- Updates of ERP software are frequent, need to be ready for potential problems

How do we respond

- From CentOS 6 to Cent OS 7
 - Tuning OS for Firebird
- From 2.5 to 3.0
 - Finally – BIG CACHE arrived!

Linux

- Linux version 3.10.0-957.21.3.el7.x86_64 (mockbuild@kbuilder.bsys.centos.org) (gcc version 4.8.5 20150623 (Red Hat 4.8.5-36) (GCC)) #1 SMP Tue Jun 18 16:35:19 UTC 2019
- CentOS 7 has much better load balancing and throughput capabilities than 6.x

Absolutely necessary Linux tuning

1. To prevent memory fragmentation and related crashes

In `sysctl.conf`

`vm.max_map_count=250000`

2. Max Open Files for 500+ (total) connections

In `firebird-superserver.service` files section
[service]

`LimitNOFile=49999`

Swapiness and memory fragmentation

Decrease swapiness (% when system starts to use swap, default is 60%)

```
vm.swappiness=10
```

memory defragmentation

```
vm.min_free_kbytes=1048576
```

not with Firebird 3!

```
#vm.dirty_ratio=60
```

```
#vm.dirty_background_ratio=2
```

Keepalive

Decrease the time default value for connections to keep alive

```
net.ipv4.tcp_keepalive_time=300
```

```
net.ipv4.tcp_keepalive_probes=5
```

```
net.ipv4.tcp_keepalive_intvl=15
```

More tuning

#Increase size of file handles and inode cache
`fs.file-max=2097152`

Increase number of incoming connections
`net.core.somaxconn=4096`

Increase number of incoming connections
`net.core.netdev_max_backlog=65536`

Increase the maximum amount of option
memory buffers

`net.core.optmem_max=25165824`

Network tuning

```
### GENERAL NETWORK SECURITY OPTIONS ###
```

```
# Number of times SYNACKs for passive TCP connection.
```

```
net.ipv4.tcp_synack_retries=2
```

```
# Allowed local port range
```

```
net.ipv4.ip_local_port_range=2000 65535
```

```
# Protect Against TCP Time-Wait
```

```
net.ipv4.tcp_rfc1337=1
```

```
# Decrease the time default value for tcp_fin_timeout connection
```

```
net.ipv4.tcp_fin_timeout=15
```

Summary for Linux

1. VMA and Max Open Files are critical for performance
2. Memory tuning of large RAM to prevent unwanted swap and other limitations
3. Optional tuning

Firebird – from 2.5 to 3.0

Fast conversion of Firebird 2.5 databases to
Firebird 3.0

<https://ib-aid.com/en/articles/fast-conversion-of-firebird-2-5-databases-to-firebird-3/>

The major warning regarding change of the SuperClassic to SuperServer

- Page cache can be set in database header
 - `gstat -h db | grep Page buffers`
- It overrides `firebird.conf` and `databases.conf`
- Very bad performance guaranteed!
- `gfix -buff 0`

```
DefaultDbCachePages = 100K      # pages
FileSystemCacheThreshold = 300K  # pages
TempBlockSize = 2M              # bytes
TempCacheLimit = 20480M         # bytes
TempCacheLogThreshold = 2046M   # HQBird only
LockMemSize = 30M              # bytes
LockHashSlots = 30011          # slots
AuthServer = Srp, Legacy_Auth
UserManager = Srp, Legacy_UserManager
WireCrypt = Enabled
WireCompression = true
RemoteServicePort = 3050
ServerMode = Super
ExtConnPoolSize = 500          # HQBird only
ExtConnPoolLifeTime = 14200    # HQBird only
SortDataStorageThreshold = 8192 #HQbird reports apply
```

firebird.conf

Databases.conf

16K pages, 220 GB

DefaultDbCachePages = 14080K

pages, up to 240 GB page cache with FS cache

FileSystemCacheThreshold = 15361K

LockHashSlots = 40099

LockMemSize = 50M

The most important settings for database

16K pages, 220 GB

DefaultDbCachePages = 14080K

pages, up to 240 GB page cache with FS cache

FileSystemCacheThreshold = 15361K

LockHashSlots = 40099

LockMemSize = 50M

TempCacheLimit = 20480M

The Games Of Cache

Total RAM, Gb	RAM for Page Buffers, Gb
314	260
314	220
314	120
314	60

Method for big databases on dedicated Linux server

1. Calculate requirements for core
 1. Normally 10-15% is fine
 2. Can be different for various Linux
2. Calculate TempCacheLimit
 1. TempSpaceCacheThreshold to identify big queries
 2. 20Gb is optimal
3. Reserve 10% for file cache
4. The rest is for Page Buffers

Why 220, not 260?

1. There are operations which require file cache
 1. Backup, Nbackup, gstat
 2. Zipping backups
 3. Service operations

Problems related with the long uptime

1. Backup speed degradation
 1. Everyday backup takes a bit longer than previous
 2. Sweep time does not change
 3. Backup goes back to normal after restart

Problems related with the long uptime

2. Time to establish connection after 1 week of time increases by 10-15%
 - Restart of Firebird helps

OS restarts

- yum-utils - needs restarting –
 - updates which require restart
- Scheduled restart once per month



PLATINUM IBSurgeon

- Replication, Recovery and Optimization for Firebird and InterBase since 2002
- Platinum Sponsor of Firebird Foundation
- Based in Moscow, Russia

www.ib-aid.com

Thank you!

- support@ib-aid.com

Linux core details

Linux version 3.10.0-957.21.3.el7.x86_64
(mockbuild@kbuilder.bsys.centos.org) (gcc version
4.8.5 20150623 (Red Hat 4.8.5-36) (GCC)) #1
SMP Tue Jun 18 16:35:19 UTC 2019