Menduliho Menduliho Millions of transactions per second on a single machine case for a virtualized database and scale-in

ROGER JOHANSSON

Who am I?

Roger Johansson

Senior Solution Architect – Starcounter http://StarCounter.io Actor Model, Scalability, Distributed Systems, C#, Go, Kotlin

Proto.Actor Founder – Ultra Fast Distributed Actors (Go, .NET, JVM) <u>http://Proto.Actor</u> <u>https://github.com/AsynkronIT</u>

Twitter: @rogeralsing Github: rogeralsing Mail: roger@starcounter.com

http://Github.com/rogeralsing/presentations

Agenda

Setting the stage - Why we need in memory computing:

- Application Platform
- Micro-Apps

The Starcounter information operating system:

- Our approach to in memory computing
- The future of hardware
- The Mars project



Front-end Framework React, Polymer

Communication Palindrom - REST, Web Sockets

Application View Models, Entities, App Logic

In Memory App Platform Mapping, Persistence, Queries

Starcounter



Micro-Applications





	Арр С

Starcounter

			🔎 Search	
Home Persons Organizations				
SAVE CANCEL				
Stefan Edqvist 🍆 🖂 ^{Starcounter AB (Sweden)}			ADD EMAIL ADDRESS ADD PHONE NUMBER	
		First name Stefan Email Primary v stefan.edqvist@starc Phone Number Mobile v 46 739856748 This person does not have any addresses yet. Click remer t	Last name Edqvist	ΰ
Organizations Starcounter AB (Sweden) MORE ORGANIZATIONS	turinday 11/11/1972 Tatalonality Sweden	Tags	Editable Custom Fields I test: x) This person does not have any custom fields yet. Click Henc to add some.	
Stefan Edqvist Username stefan FirstName Stefan LestName Edqvist				
Authentication settings servew PASSWORD RESET PASSWORD Member in				
Name Admin (System Users)			Action	
dd to a system user group Vone			(+ ADD TO THE GROUP
DELETE UNDO GLOSE	Title Wints	e te Note here	SAVE	

	🔎 Search	
Home Persons Organizations		
SAVE CANCEL		
Stefan Edqvist 📞 🖂	ADD EMAIL ADDRESS ADD PHONE NUMBER	CREATE ALL 0
	First name Last name Stefan Edqvist	
Organizations Birthday Starcounter AB (Sweden) 11/11/1972 NoRE ORGANEZMONS NoRE ORGANEZMONS NoRE ORGANEZMONS NoRE ORGANEZMONS Nore Nore Nore Nore Nore Nore Nore Nore	Tags Editable Outrom Fields CB Net Core Itest This person does not have any outrom fields yet. Click HERE to add some.	
Stefan Edqvist Username stefan FirstName Stefan LastName Edqvist		
Authentication settings		
Name	Action	
Admin (System Users) Add to a system user group None	REMOVE	AGO TO THE GROUP
DELETE UNDO CLOSE		
	Title Write Note here C SAVE PRIVACY TOGGLE GO TO OWN NOTEBOARD test This is a note	





Data Flow Programming

SUM	• I	$\times \checkmark f_x$	=B2*B1
	А	В	С
1	Price	5	
2	Quantity	10	
3	Total	=B2*B1	
4			



Demo Dataflow

In Memory, For Real



Database Process



Database Process





C# Source Code

```
public class Person
{
  public string FirstName
  {
    get;
    set;
  }
```

C# Compiler Output

```
public class Person
{
 private string _firstName;
  public string FirstName
   get => _firstName;
    set value => _firstName = value;
  }
```

Starcounter Weaver Output

```
public class Person
 private string firstName;
 private ulong id;
 public string FirstName
   get => DbState.GetString(_id, FirstNameId);
   set value => DbState.SetString(_id, FirstNameId, value)
```

Millions of operations per second

O/R mapper – Object Graphs



Lazy Load Execute SQL Query per traversed relation

Eager Load Carefully defining load boundaries Ripple load, Hundreds/Thousands of queries per view

Hard to maintain, inefficient joined queries

Starcounter – Object Graphs



Chasing Pointers

Traversing relations are simple pointer traversals



 $\, \approx \, \, \lor \,$

😌 Market Input			E) ($\Theta \Theta \Theta$	€ €	0 0	Θ						Module Variant
⊗ Engineering Input										1				MODULE VARIANTS
𝔝 Module Concepts														mover failants)
♥ Product Configuration				Цe,										
𝔝 Profitability				Ded				an		a la				
le Supply Chain				t Pro	<u>i</u>			e text	inlet	deys	2			
scs Supply Chain Structure				po	SeP	pue	gion	rfac	jo aj	yling	Ē			
VP Volume Planner				۱.	å	<u>ه</u>	8	25	5	5	A			
🐟 Navigator				11	USD	-	-	VDI	mm	-	I/min			
(march)				S	s	s	s	s	s	s	s			
i an an an an an				C	в	V V	v	в	V	E	E			
P 8	(1)			#	1	5	3	1	2	2	6			
ji cov c.v	ŏ	Modules / Variants	Description									VCP		
· ···· ···· ···· ··· ··· ··· ··· ··· ·	\odot	 Shell module 	Shell and container Styling							•		6	A	
THE REPORT OF THE PARTY AND A	(\downarrow)	Shell, Extended	Extended, w "dust wheel"							Base style				
Nac	õ	Shell, Short	Short, Compact							Exclusive				
ji war	Φ	Shell, Long	Long, Standard							Base style				
		 Dust flap module 	Dust flap One-way-valve								0	1		
	· 🕞	Dust flap, long life	Long Life, Silicone								120;1			
	(+)	Dust flap, std	STD, Black rubber								80;90;			
	ŏ	Charger connect m	Charge plugg and Charge so									1		
	0	Dock module	Wall mounting, Docking Uni							•		2.25		
		▶ PCB module	Printed circut board, electro									1		
	\bigcirc	Speciality Nozzle m	Crevice nozzle (detachable)	1					•		0	4.5		
	0	Nozzle module	Nozzle (fixed)	1		•	0	0	•		0	7.5		
	\cup	Control module	On/Off switch, Container rel					0				1.5	1	
		Container module	Container for dust collection					0				3		
		Handle module	Handle, used for Styling					0				1.5	11	
		Charger module	Power Adopter (transformer									54	11	
		Base module	Shell base									1		
		Filter module	Dust filter(Detachable)									3		
		Container styling m	/					0		•		6	+	
										•				
											•			

Demo - Write 1 mil transactions

Demo - Write 1 mil transactions

In this demo we provided:

- ACID Transactions
- Snapshot Isolation
- Guaranteed disk persistency

Can we go even faster?

- We are capped by disk IO
- The Mars project

The Future of Hardware



Price of 1MB in USD over time

Registers

Cache

Main Memory

Solid State Disk

Magnetic Disk

Cache Main Memory Non Volatile Memory Solid State Disk

Registers

Magnetic Disk

Current

Future

Starcounter on the CPUs of 2018



Few Big Servers, not many small

Few Big Servers, not many small

QEA

Starcounter.10