

Patterns for scalability and availability in (trading) systems

Michel André – CTO – Saxo Bank



Saxo Bank – introduction

- Global online investment bank – facilitator/broker setup - offices in 25 countries clients in 180 countries
- Specialises in online trading and investment, servicing retail clients, corporations and financial institutions
- A leading presence in online trading due to client service, competitive pricing and industry-leading trading platforms.
- Enables private investors and institutional clients to trade FX, CFDs, ETFs, Stocks, Futures, Options and other derivatives via multi-award winning online trading platform.
- 3rd generation technical platform and evolving – Microsoft based, mostly custom developed in house
- 10 of thousands of concurrent users, 100 of thousands of price updates/sec, very high transaction peaks around numbers and market state changes.



Core Business Model



Cloud based – Platform as service – White label

Truly multi-asset / single account

OTC

Forex



- ✓ 160+ crosses
- ✓ Tradable quotes
- ✓ Request for quotes
- ✓ Margin

CFD



- ✓ 8700+ stocks, 22 index trackers, 20 commodities
- ✓ Tradable quotes, extensive liquidity
- ✓ Algorithmic orders/Smart order routing/DMA

FX Options



- ✓ 40+ crosses
- ✓ Tradable quotes
- ✓ Broad coverage 1d-1y
- ✓ Vanilla, binary touch

Bonds



- ✓ Wide range sovereign, government, corporate bonds
- ✓ Offline traded
- ✓ Collateral

Exchange Traded

Stocks



- ✓ 18400+ equities
- ✓ 33 exchanges
- ✓ Algorithmic orders/Smart order routing/DMA

Futures



- ✓ 230+ base contracts
- ✓ 22 exchanges
- ✓ Margin

ETFs & ETCs



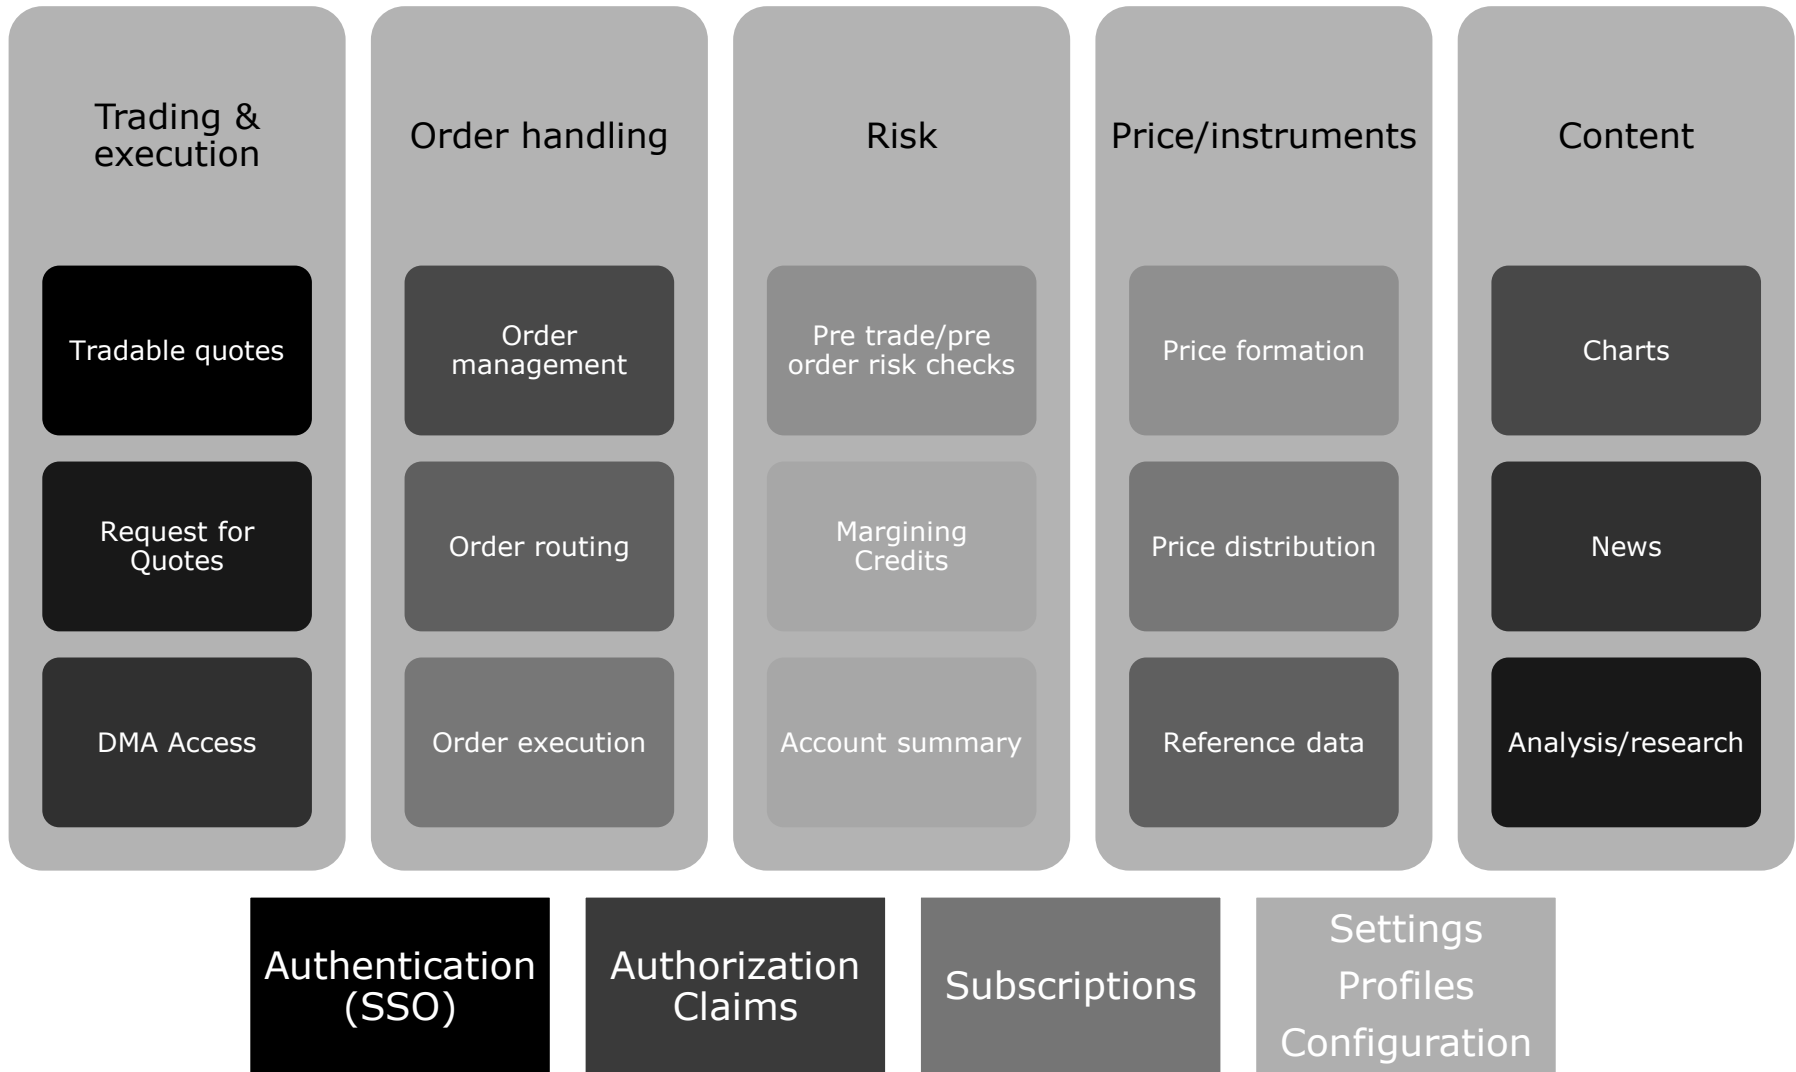
- ✓ 2270+ exchange traded funds (ETF), commodities (ETC)
- ✓ Listed, traded and settled as stocks

Contract options



- ✓ 66+ base contracts (stock indices, commodities, interest rates, forex)
- ✓ 17 exchanges
- ✓ Margin (SPAN/Portfolio)

Workflow/service requirements



The numbers...

Over **13,000** concurrent online clients - operational and open 5.5 days * 24 hours

In excess of **400,000** price feed updates per second

Intraday execution of up to **900** trades per second

Over **500,000** trades booked and processed daily

160,000 unique users monthly to our trading portal

While maintain latencies in the single digit millisecond range throughout

Business challenges

Unified and compelling client experience

- Across devices/platforms
- Across products
- Across segments
- Across geographies

Cost-efficiency

- Technology sharing
- Reuse
- Avoid duplication and proliferation

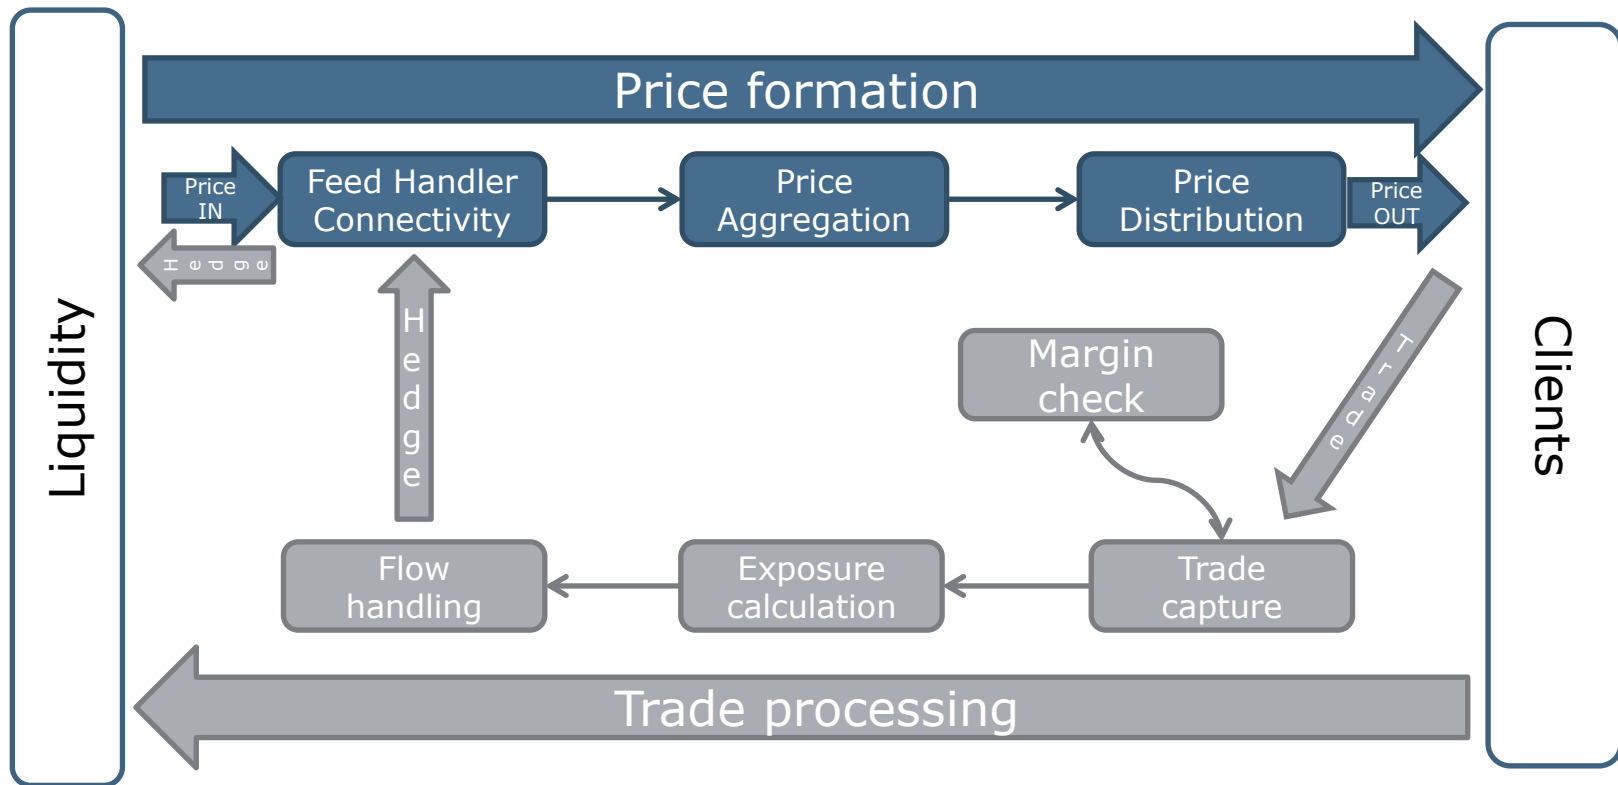
Faster time to market and flexibility

- New products
- Features/sophistication
- New geographies

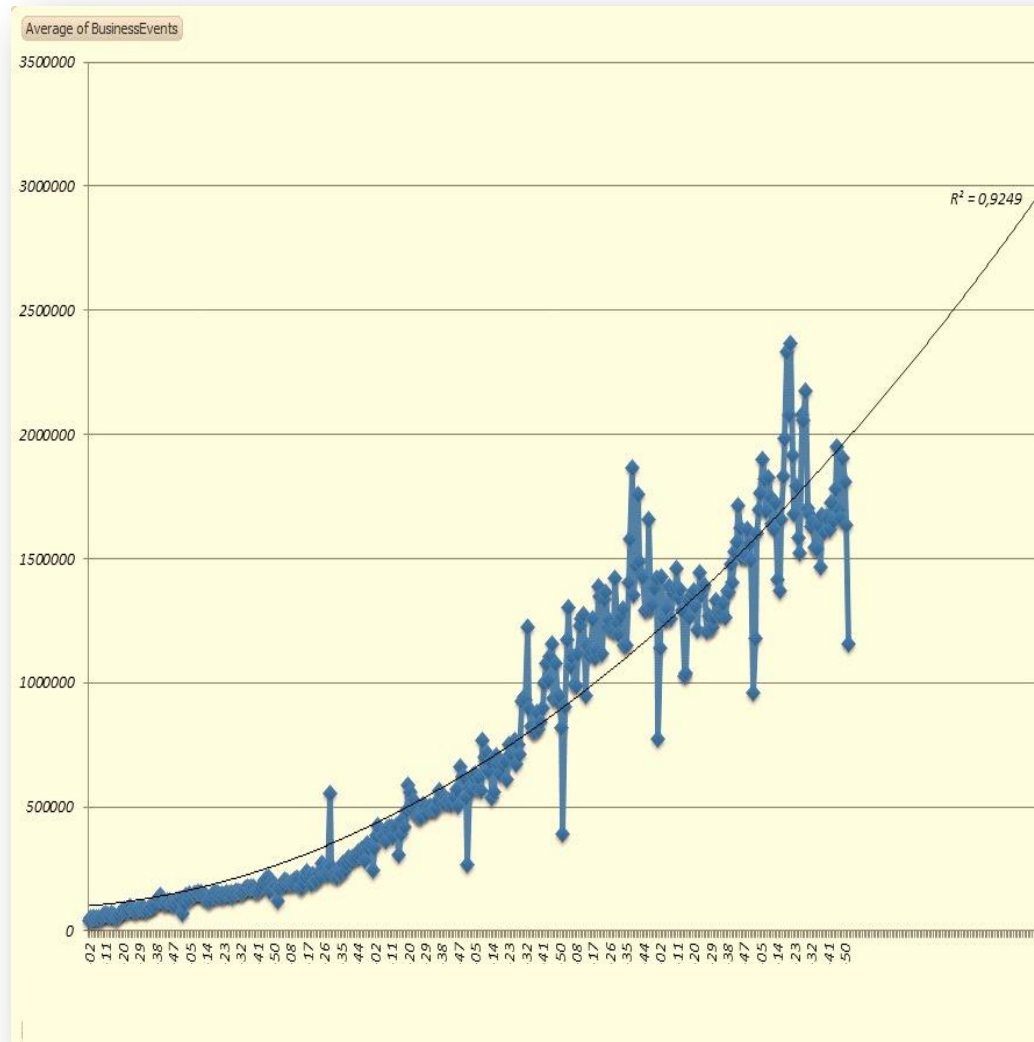
Compliance and regulation

- Regulated in many geographies direct/indirect

Critical business flows – where milliseconds matter



Challenge - Business growth – with corresponding data growth



Trading systems are "Reactive systems" - at heart ...

...core design principles mapped to tenets of reactive

Responsive

- Casual/eventual consistency guarantee
- Horizontal scaling/partitioning of work sets
- Asynchronous/fire and forget interaction where it can be supported
- Very high throughput and tight latency requirements

Elastic

- Scale out on commodity hardware
- Horizontal scaling/partitioning of work sets
- Swift capacity growth by horizontal scaling

Resilient

- Redundancy (active/active)
- No single points of failure
- Automated seamless failover
- Survive with reduced service a datacenter failure/outage
- State replication

Message Driven

- Event driven / Message based
- Loosely coupled through topic based pub/sub (location independent)

<http://www.reactivemanifesto.org/>

Technical Environment

Microsoft through and through

- Windows 2008/2012 servers
- SQL Server 2012

Middleware/Messaging

- 29West / Informatica (Multicast messaging and persistence)
- WCF
- MS Biztalk (Integration backoffice/enterprise)
- Plain old sockets

Standard Windows Services (.NET C#, C++, IIS 7.5)

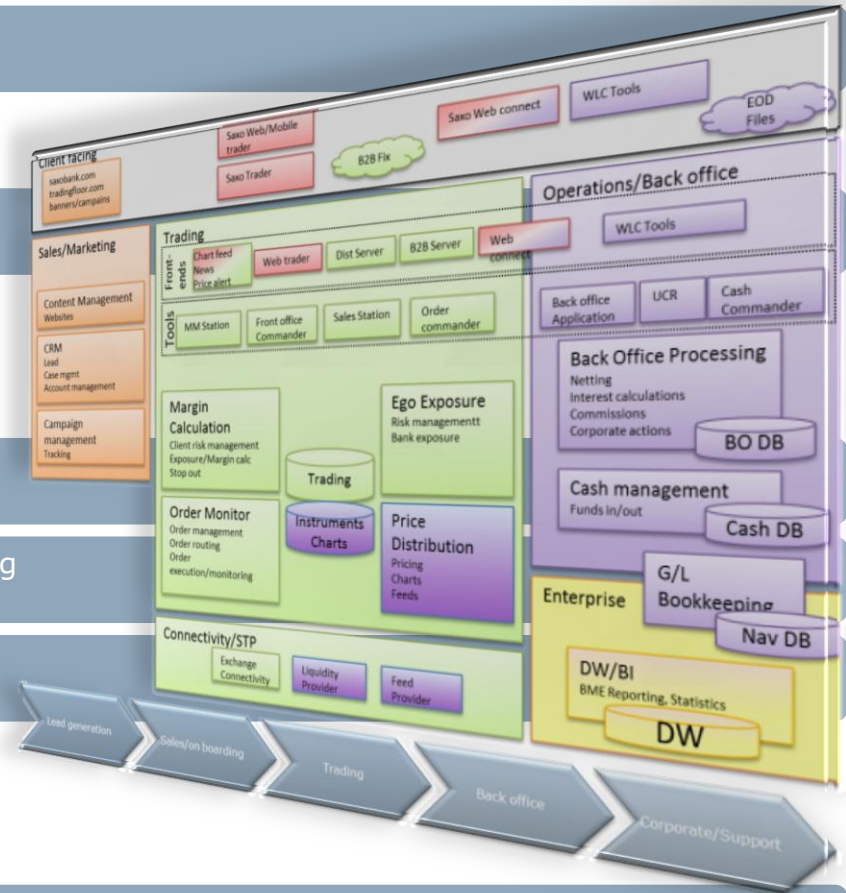
Shared in process components for business logic and caching

Monitoring, instrumentation & tracing

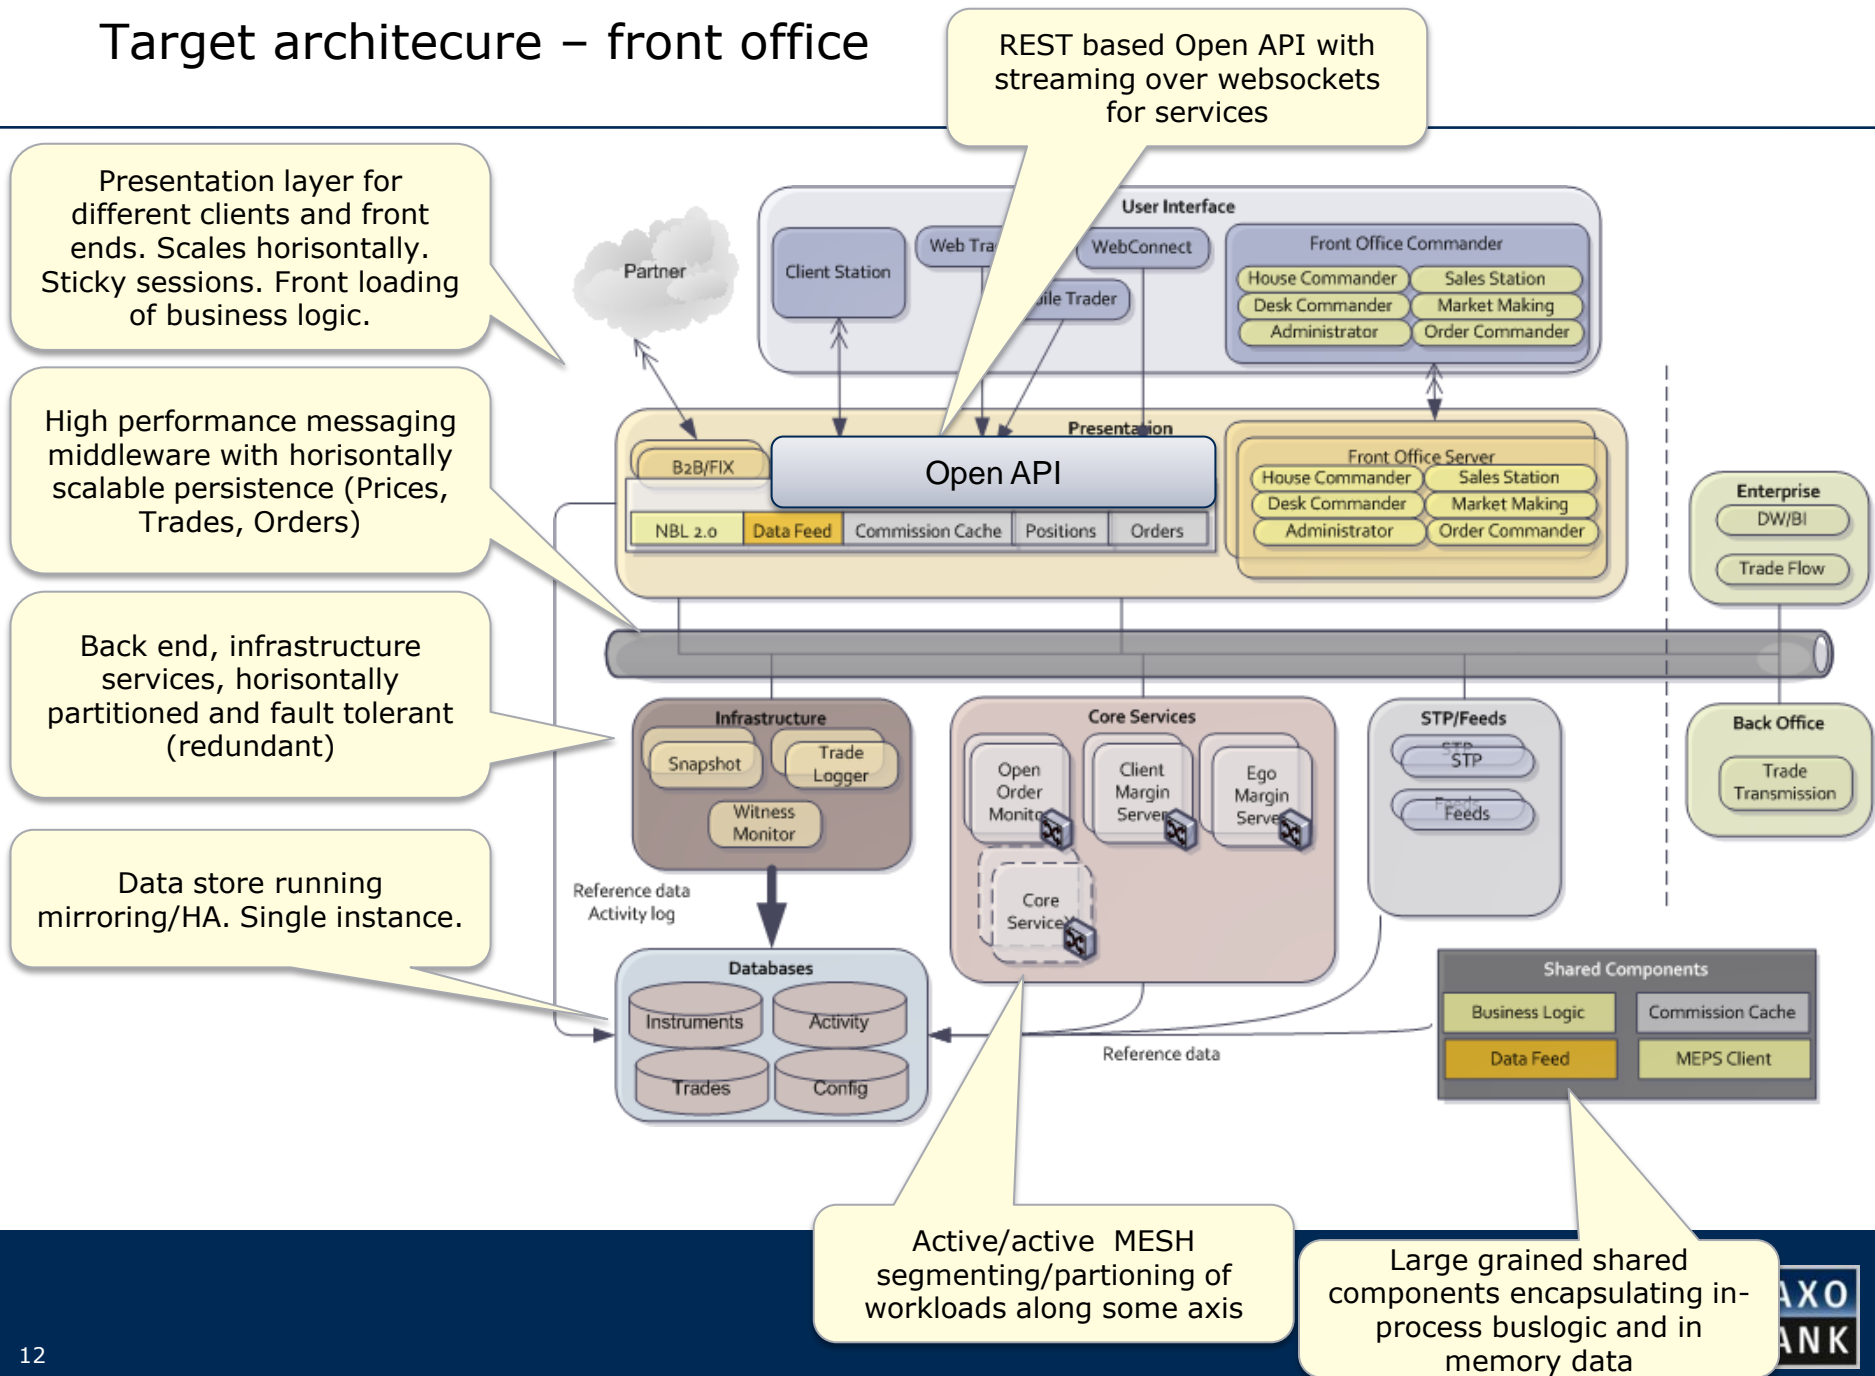
- Velocimetrics
- SCOM
- Standard Windows performance counters
- Eventlog and Logfiles

Cisco infrastructure

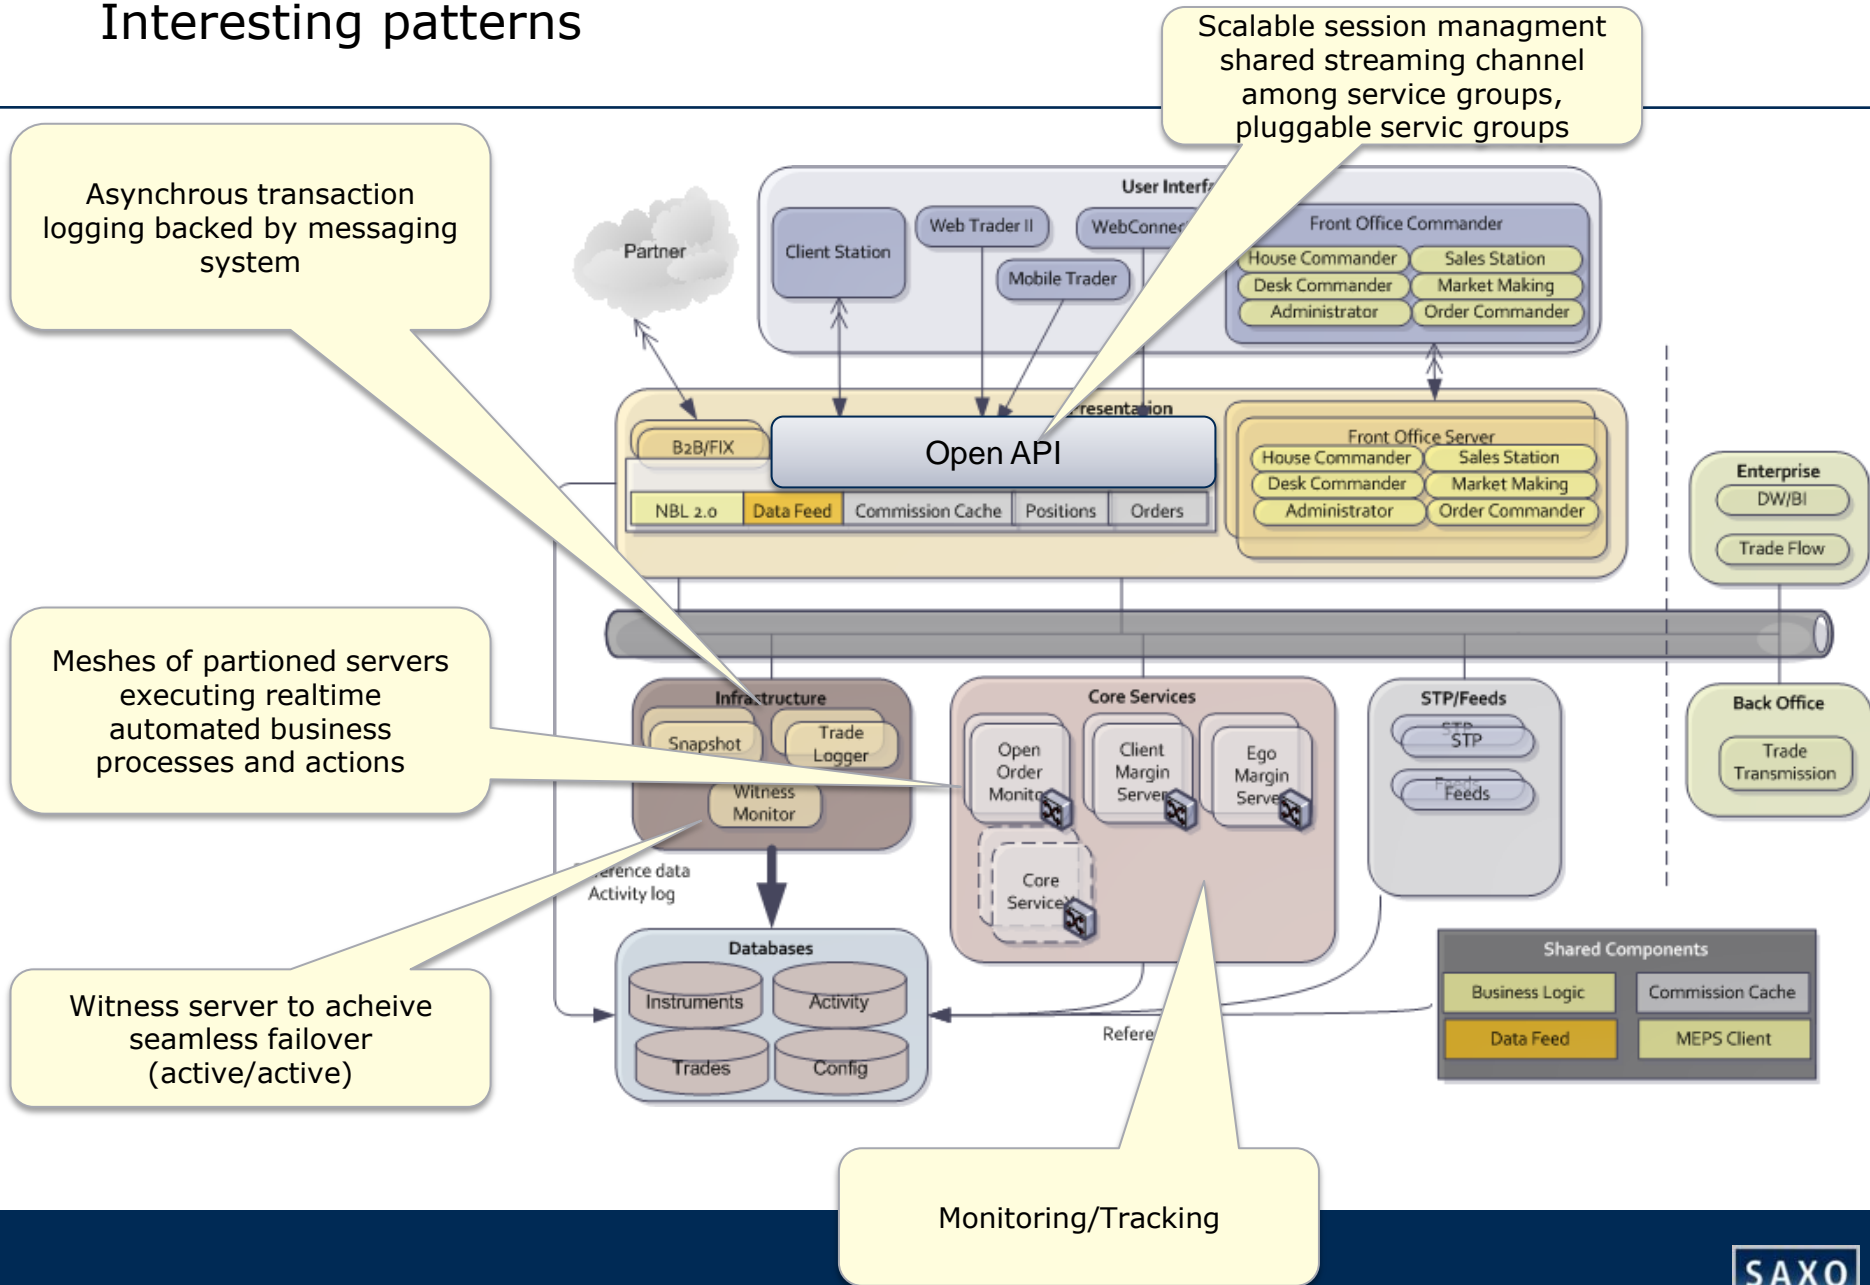
- ISP Grade
- MPLS



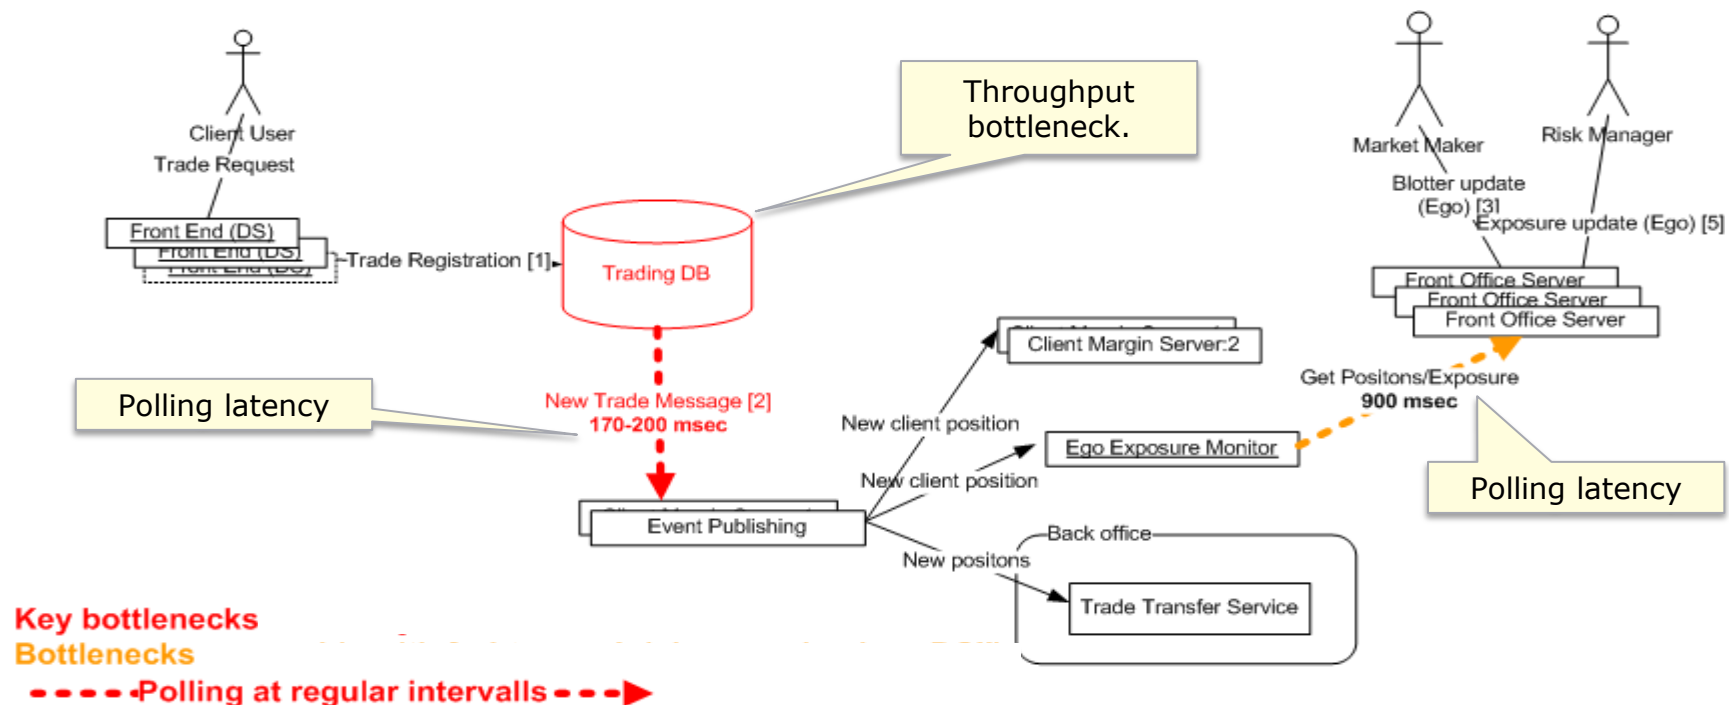
Target architecture – front office



Interesting patterns

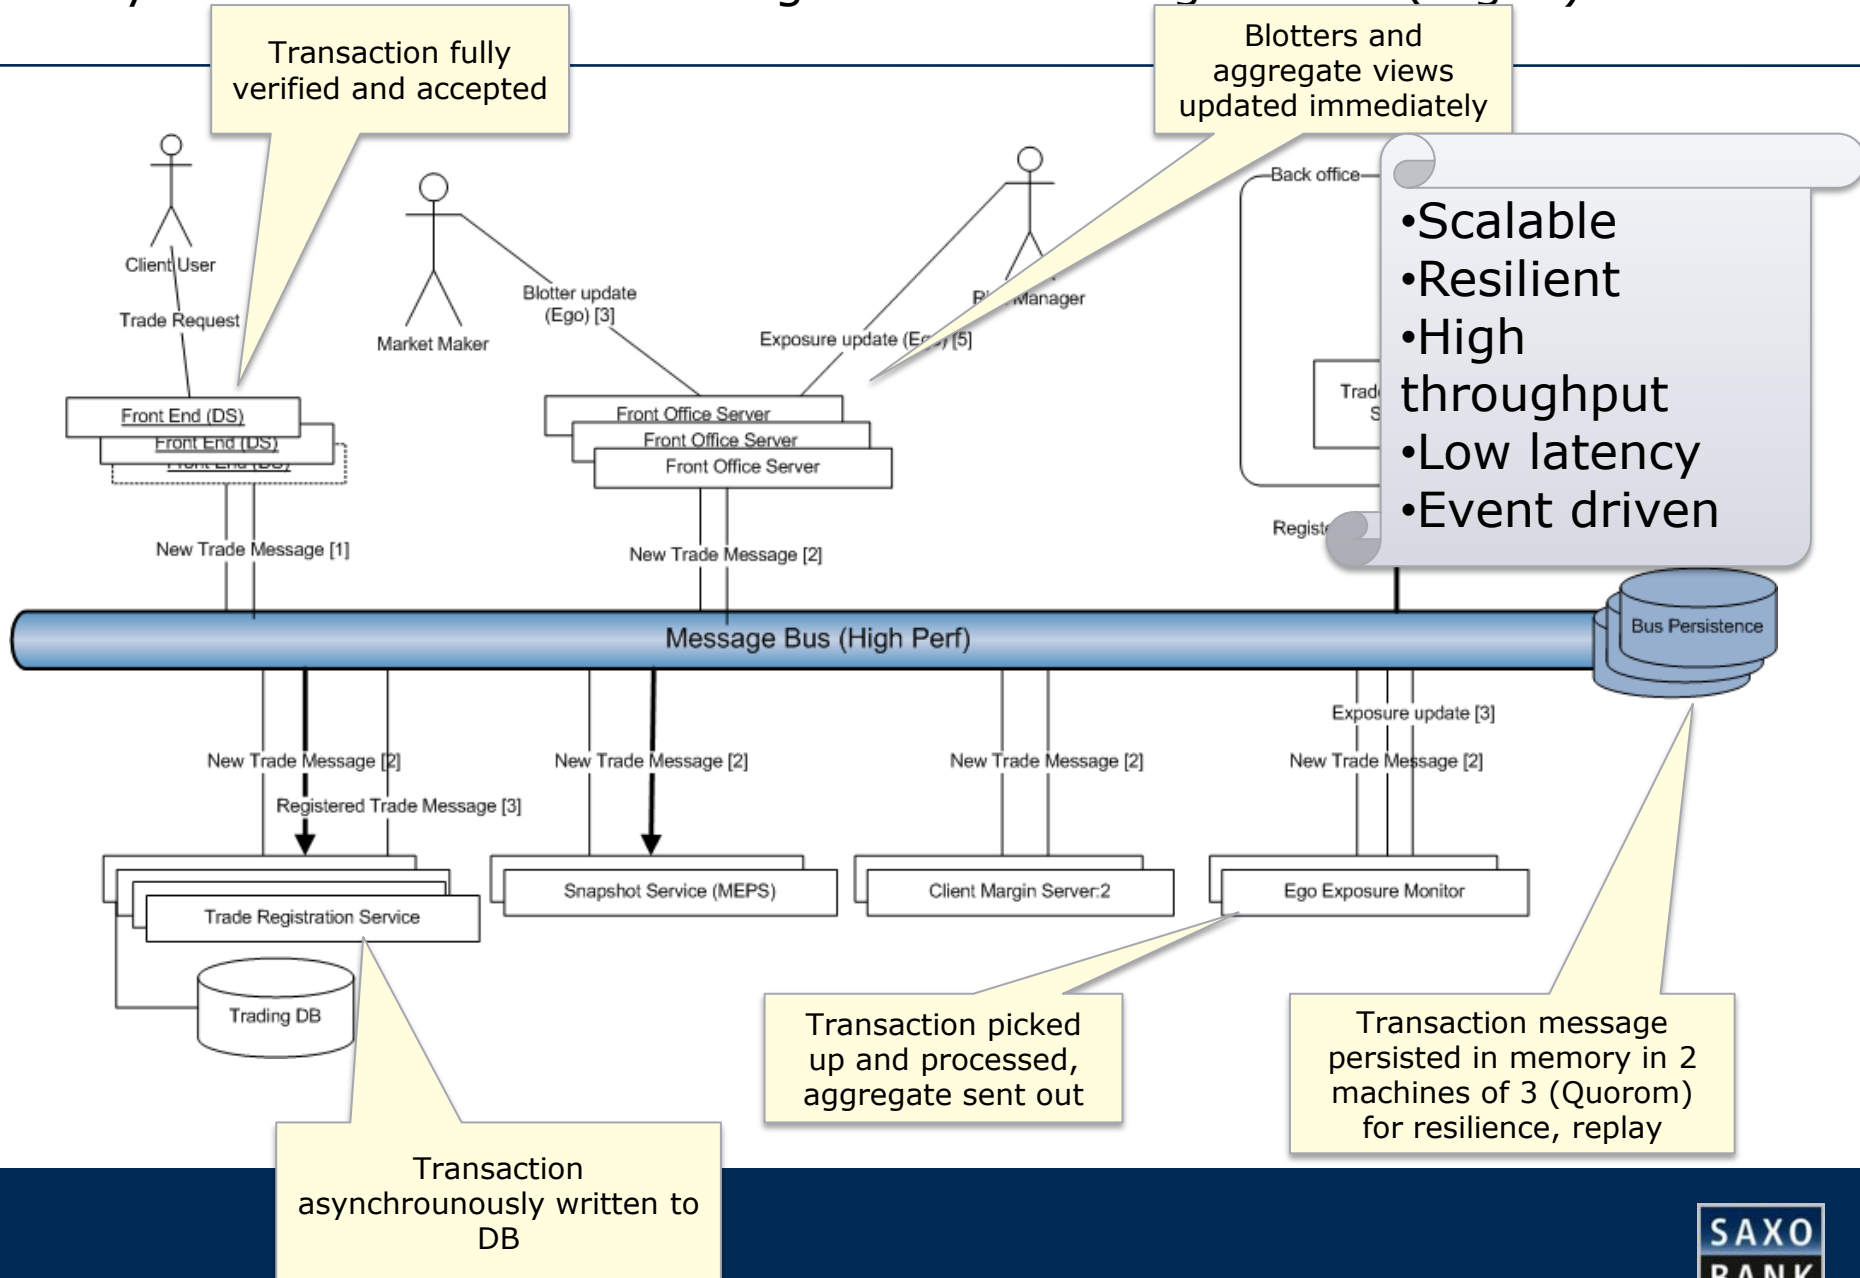


Traditional architecture – start state

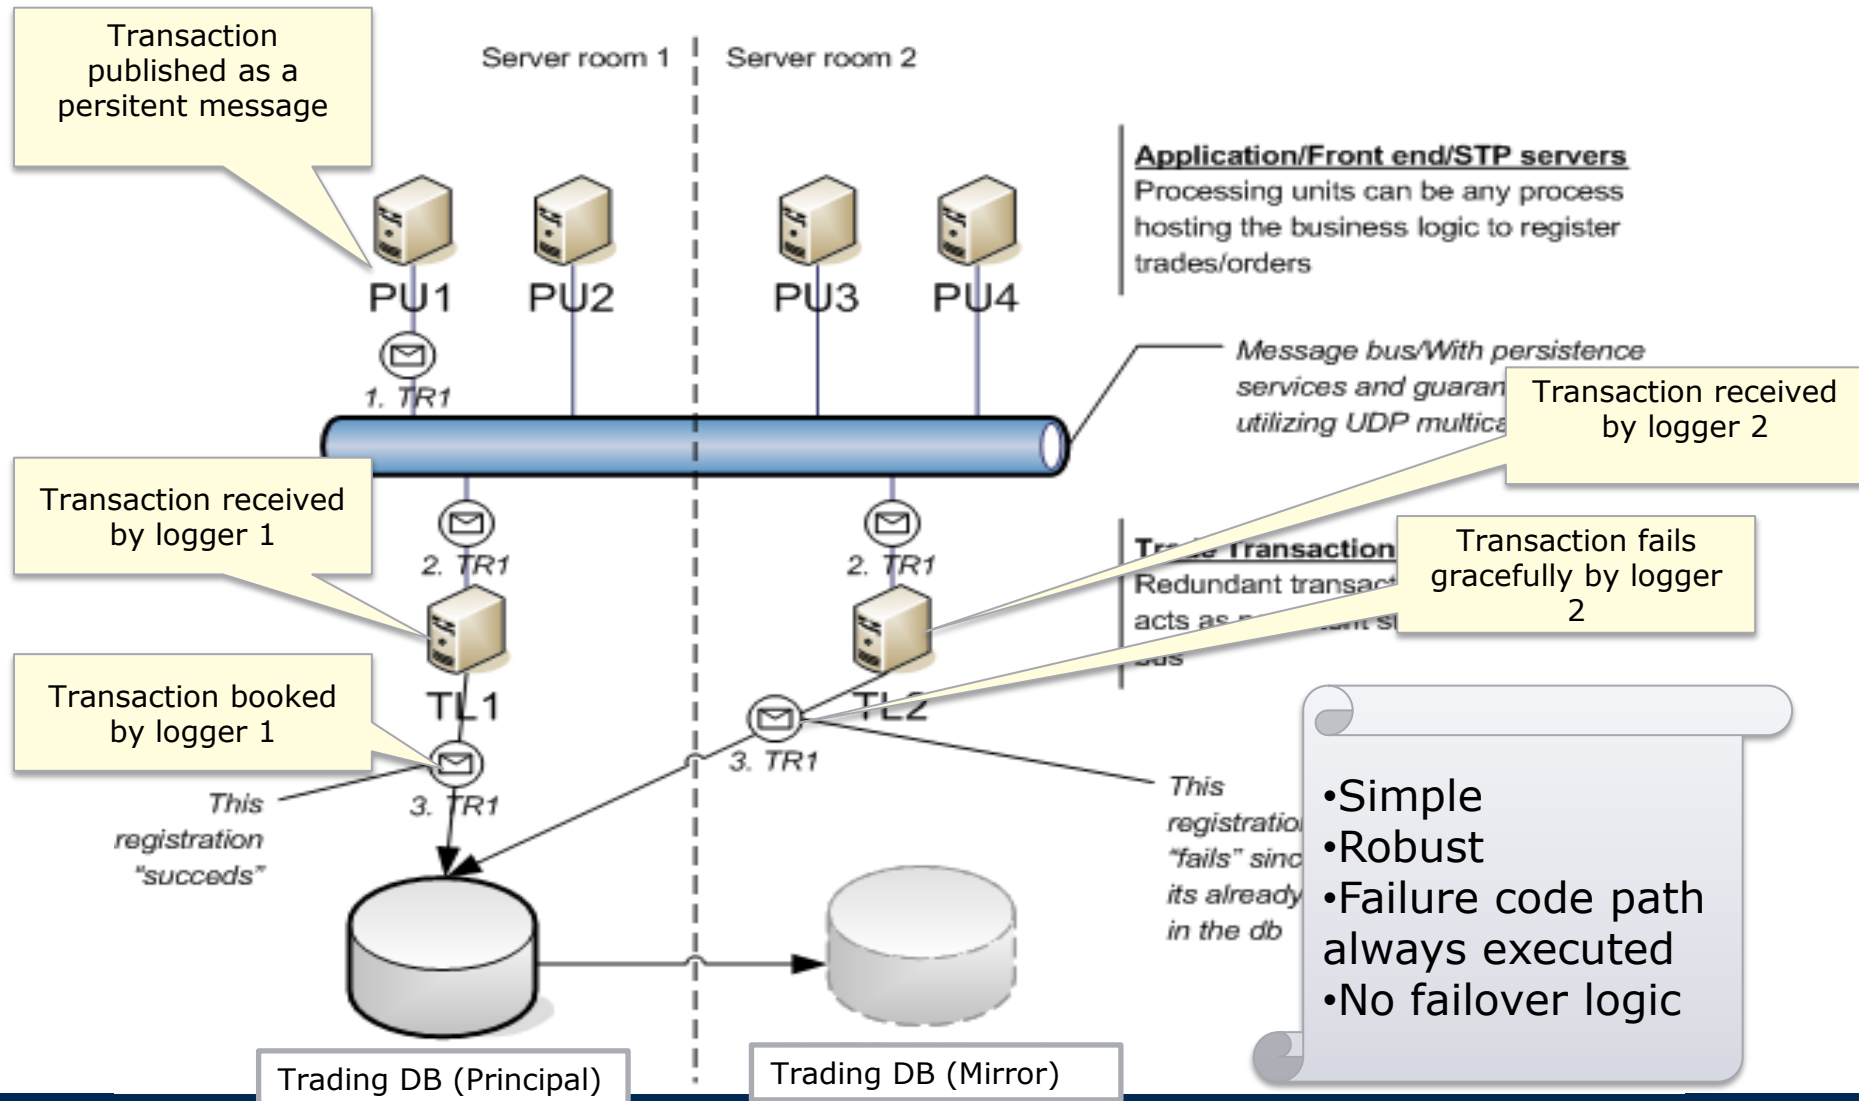


- Long latency
- Database bottleneck
- Capped throughput (by DB)

Asynchronous Transaction Registrations – Target state (3 gen)



Asynchronous Transaction Registrations – Final Twist (fully active)



It's not all architecture and technology



Learning points - liabilities

- Architecture needs resilience, scalability, monitoring, supportability built in and taken care from the start due to complexity and number of nodes/ci
- Synchronize projects and implementation with infrastructure (machine rooms, firewalls, network segmentation - is also part of the solution)
- Do gradual rollout and provide rollback and backwards compatibility to decrease operational risk of large architectural impacting changes.
 - This comes at a cost both in development and phasing out of old style clients
- Test and verify thoroughly
- Your architects might get skinny and loose some hair

