

Extreme Performance and the Cloud

Mark Campillo – SVP Product Engineering and CTO of NWEA



Coming Attractions ... (Agenda)

- Who is NWEA?
- The Challenge
- Our Approach to Fix
- The Storage Subsystem
- What We Learned
- Q&A




Who is NWEA?

- Northwest Evaluation Association (NWEA)
- Global non-profit founded in 1977 in Portland, OR
- Founded by Educators and Assessment Researchers
- Mission: “**Partnering to help all kids learn[®]**”
- Serving more than 8 million kids and more than 25 million education stakeholders in 120 countries
- Decades of Historical Assessment Data

The Challenge

- Primary Product Project Failure
- Primary Product Failing to Perform and Scale
- High Stakes Outcomes
- Vocal and Passionate Customers
- Usage Model Perfect for the Cloud
- Built for Other than the Cloud
- Failing < 500 Concurrent Users



The Plan and Target

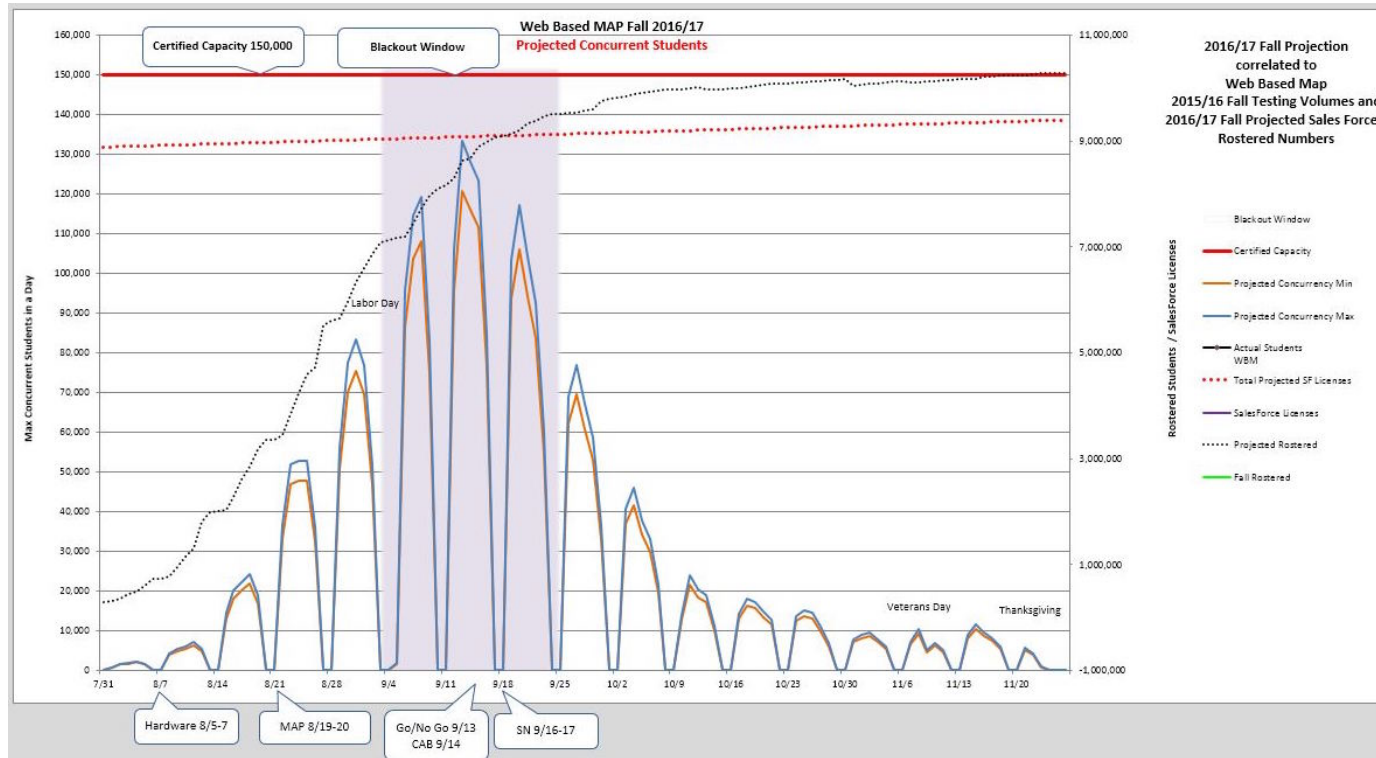
1. Stop the Leaks, Right the Ship
2. Apologize to Customers
3. Re-Design and Re-Implement
4. Scale 10x in 12 months
5. Sub-Second over the Internet(!)
6. Take it to the Cloud (2011 - 2012)
7. Retire to Paradise



The Beginning, Middle and End



Our Usage Model and Pattern



The Approach Abridged (2012)

1. Build Performance Test Suite on Cloud
2. Test New Architecture on Cloud
3. Build Prod on Cloud
4. Tune, Buy, Tune, Build, Tune, Configure, etc.
5. Go Bigger, Go Faster!
6. EOL our internal Infrastructure

The Hosting Model

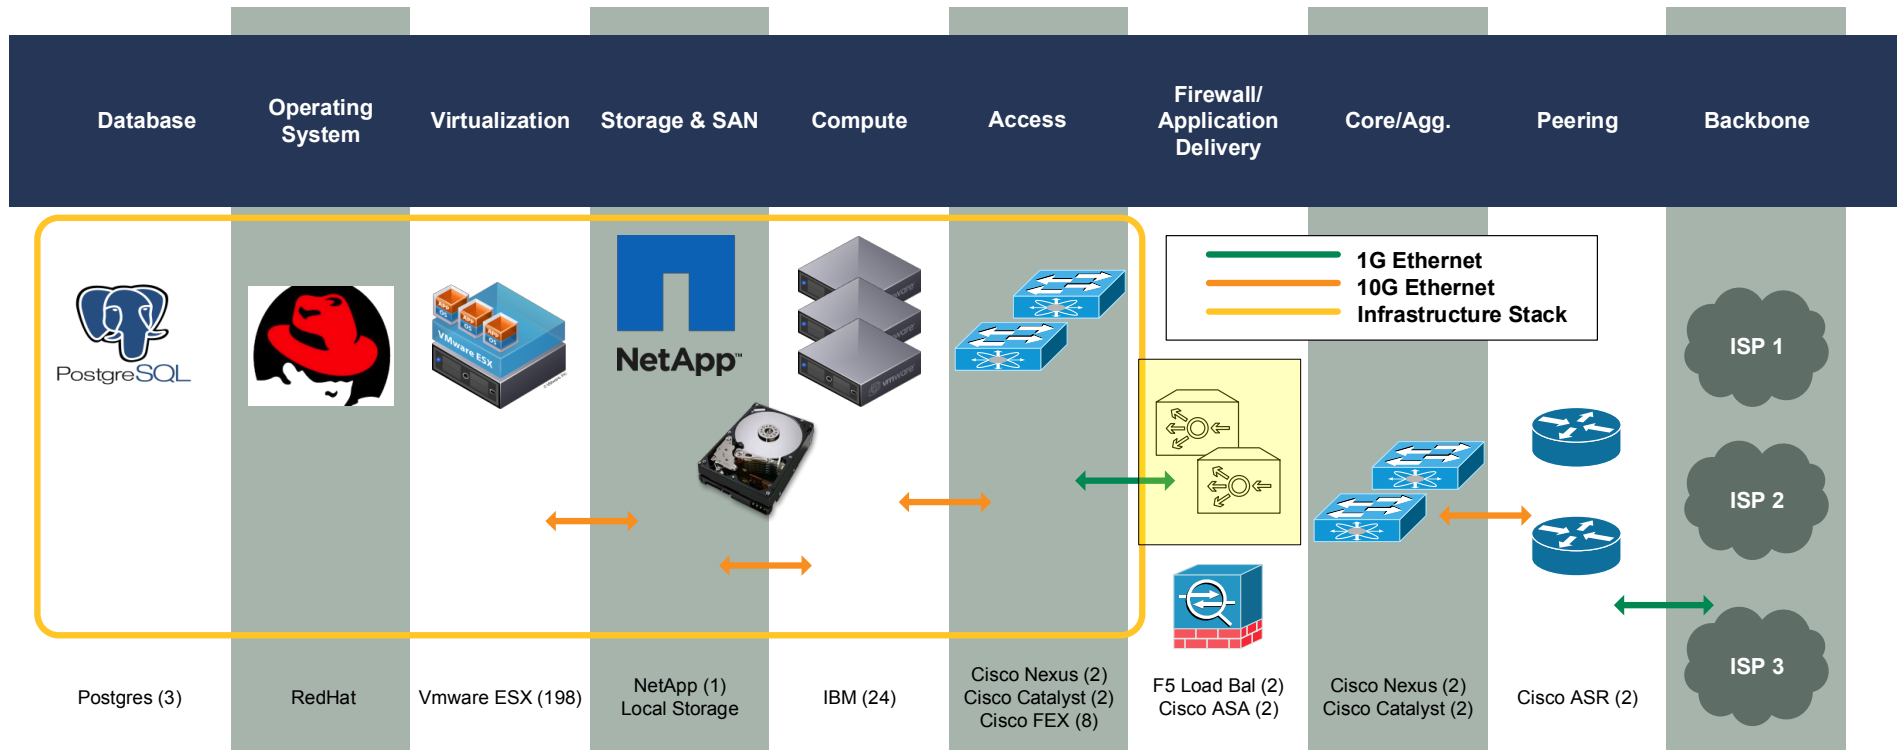
The Start

- Legacy Data Center
- Static DR (Minimal)
- No Performance Test Environment
- No Cloud Services
- Start on Virtualization

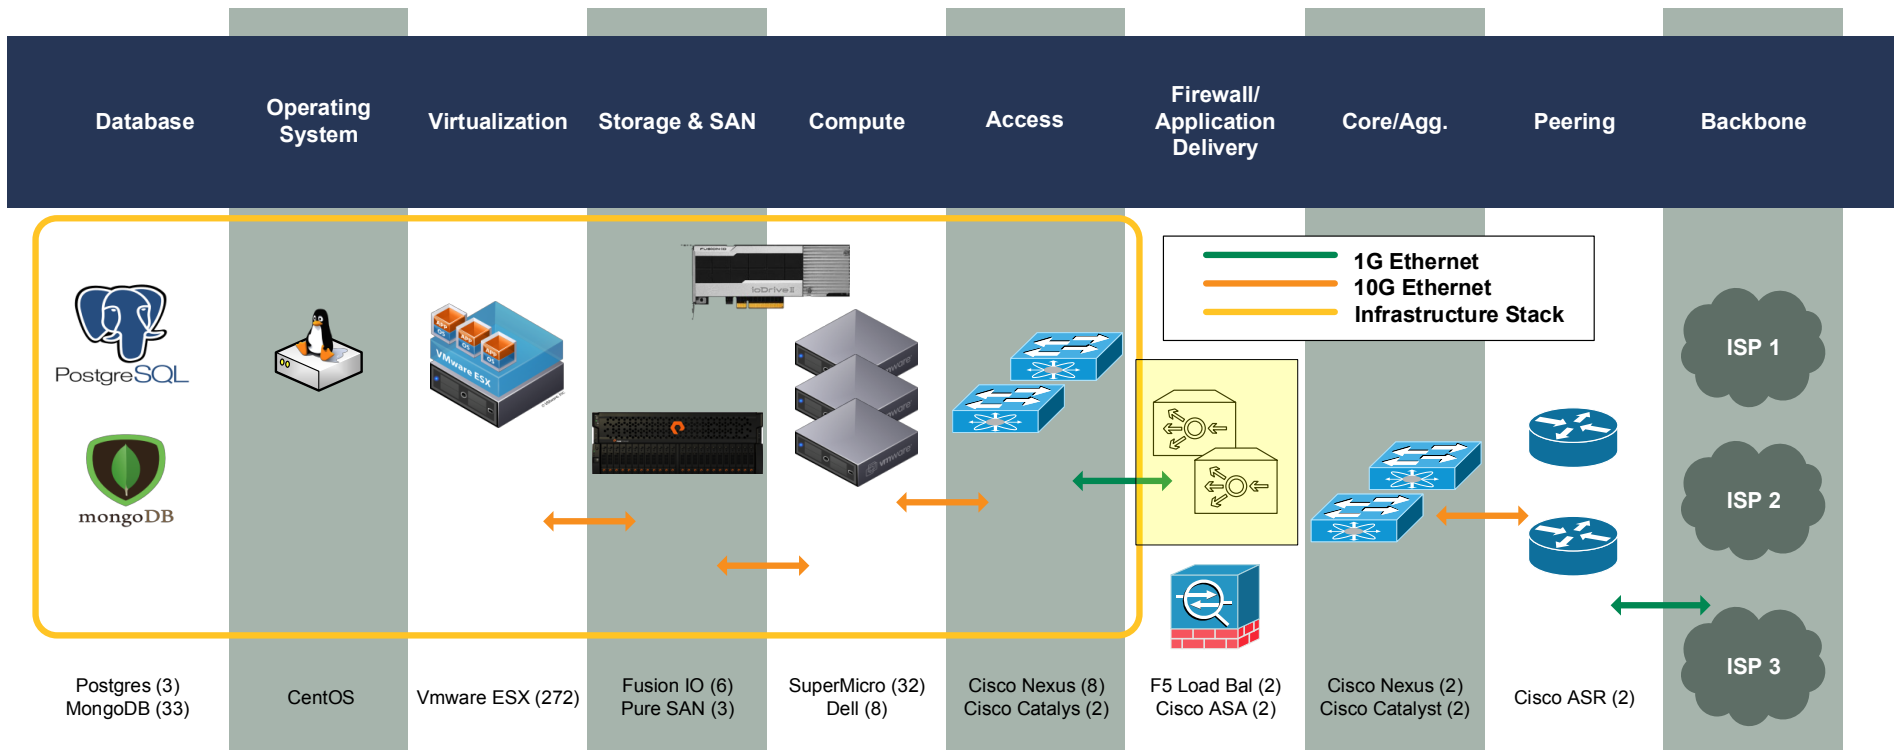
Current State

- New Tier 3 Data Center
- DR and Performance Environments
- Cloud for Load Generation
- Newest Product on Cloud
- Services Architecture
- Private Cloud prep for Public

The Stack - Start



The Stack - End




The Storage Sub-System

One of Many Optimization Points




Storage Sub-System – Starting Point

- Distributed Servers with Local Storage
- Spindles and such ...




Then came the SAN

- Spinning Disk SAN
- An addition to the local storage system



It Didn't Work! What now?

- Our Legacy SAN Sucks! (KIA for Fusion IO)
- Maximum reduction in Latency and increase in throughput
 - Fusion IO
 - Pure Storage



Some Numbers on Performance and Capacity

	SAN (Spinning)	Fusion IO	Pure Storage
Write Latency	> 100 ms	< 70 μ s	< 1 ms
Write IOPS	2,500 – 4,000	> 450,000 (4K)	> 200,000 (32K)
Bandwidth	1.5 GB/s	2.5 GB/s	> 6 GB/s
Our Scale	49K	125K	125K
Response Time	> 250 ms	< 25 ms	< 25 ms

Platform	Concurrency	Tests	Data Growth	MB/Test
Old Gen	29,660	5,734,243	2.75 GB	0.51
Next Gen	82,803	13,004,758	26.38 GB	2.41

Pros and Cons - SAN

- Pros
 - Cost of Entry
 - Scalability
 - Ubiquitous Support
- Cons
 - High Latency
 - Low IOPS
 - Low Bandwidth
- Our Result
 - Sub-Optimal Performance
 - Frustrated Users
 - Frustrated Infrastructure Engineers

Pros and Cons – Fusion IO

- Pros
 - Low Latency (almost none)
 - High IOPS
 - Good Bandwidth
 - Tested at Scale (by us)
- Cons
 - High Cost
- Our Result
 - Optimal Performance
 - Good User Satisfaction
 - Frustrated Infrastructure Engineers

Pros and Cons – Pure

- Pros
 - Low Latency
 - High IOPS
 - Excellent Bandwidth
 - Scalable in Theory
- Cons
 - High Cost
 - Untested at Scale (by us)
- Our Result
 - Optimal Performance
 - Good User Satisfaction
 - Satisfied Infrastructure Engineers
 - Ideal DR Scenario

What DID NOT Work

1. Build on What We Had
2. Tune the RDBMS
3. Tune the Application
4. Tune the Infrastructure
5. Upgrade for Capacity – i.e. Buy MORE
6. View the Cloud as a Silver Bullet Solution

Result: Marginal or No Gains – (i.e. Not Good Enough)

What DID Work

1. Quickly Adapt to Learning
2. Throw Out What We Had
3. Throw Out the RDBMS for NoSQL
4. Re-Write the Application (Micro-Services)
5. Replace Storage Architecture (Fusion IO)
6. Make Reliability the Focus
7. Build a Private Cloud


Result: >150x Scale – Removed ALL primary failure Modes

Move to the Cloud? (We seem a perfect fit)

- Our model did not perform (2012)
- We now have a major success with internal infrastructure
- We are recovering from torturing our customers
- We have a huge feature backlog
- We need to be convinced that the IOPs on the DB is up to scale
- Reliability is a Competitive Advantage
- SaaS Assessment Platform (>130K RPS, >100TB and Growing Fast)
- Can it be done on the Cloud?

What is Next for Us?

1. Major Feature Builds
2. Usability Upgrades
3. No Static DR – Load Balance Sites
4. Get out of the Datacenter Business
5. Move to Utility Computing (Cloud)
6. Retire to Paradise



Thank You!

- Mark Campillo
- SVP Product Engineering and CTO
- LinkedIn: [linkedin.com/in/markcampillo](https://www.linkedin.com/in/markcampillo)
- Email: mark.campillo@nwea.org

