Building a Hybrid Cloud at Canadian Pacific

Stuart Charlton,
Director –
Infrastructure & Operations
Information Technology

CANADIAN PACIFIC
DRIVING THE DIGITAL RAILWAY »

Canadian Pacific in 2010

14,800

mile network

15,500

active employees

$5.0

billion in revenues

77.6

operating ratio

Canadian Pacific’s Network

Vision: To be the safest, most fluid railway in North America

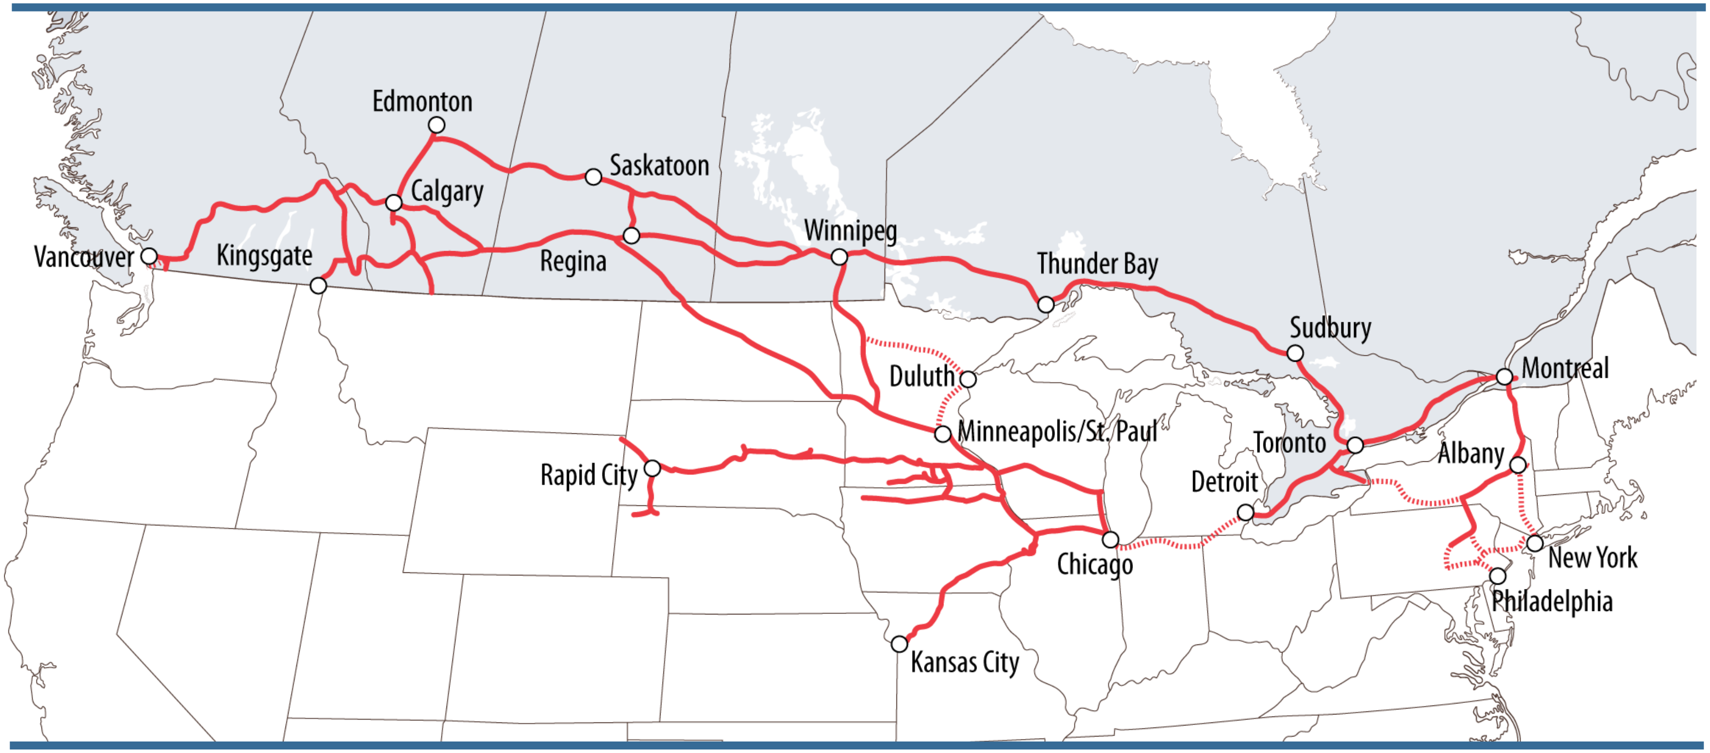

CP operates in 6 Canadian provinces and 13 US States

IT Transformation 2009-2015

Responding to the Railway Industry's Global Renaissance…

- Integrated Information Program
  - First Joint IT/Business Strategy
  - Big SAP Investment
  - Big Legacy Revitalization
- Positive Train Control
  - Integrated C&C
- Predictive Operations
- New Ordering Processes
  - Canadian Grain
- Reducing Operating Ratio

- Givens:
  - Major IT capital reinvestment starting in 2010 (more than doubled)
  - Planned for IT to deliver more in a single year than was done in prior 8 years combined

Our Assumptions

- Challenge #1: Volume, lead times & costs of **infrastructure**
  - Timeframe: 2010+

- Challenge #2: Bending down the **operational** cost curve for production
  - Timeframe: 2011+

- Challenge #3: Reducing ***cycle time*** of delivering changes to systems
  - Timeframe: Pilot 2011, Rollout 2012+

- Challenge #4: Increasing the **availability** of core operational systems
  - Timeframe: 2012+

Approach: Using the right tool for the job, given the time constraints

Caveat: Forward-looking - this all may change

Advice we got: "Look at how complicated all this stuff is!"

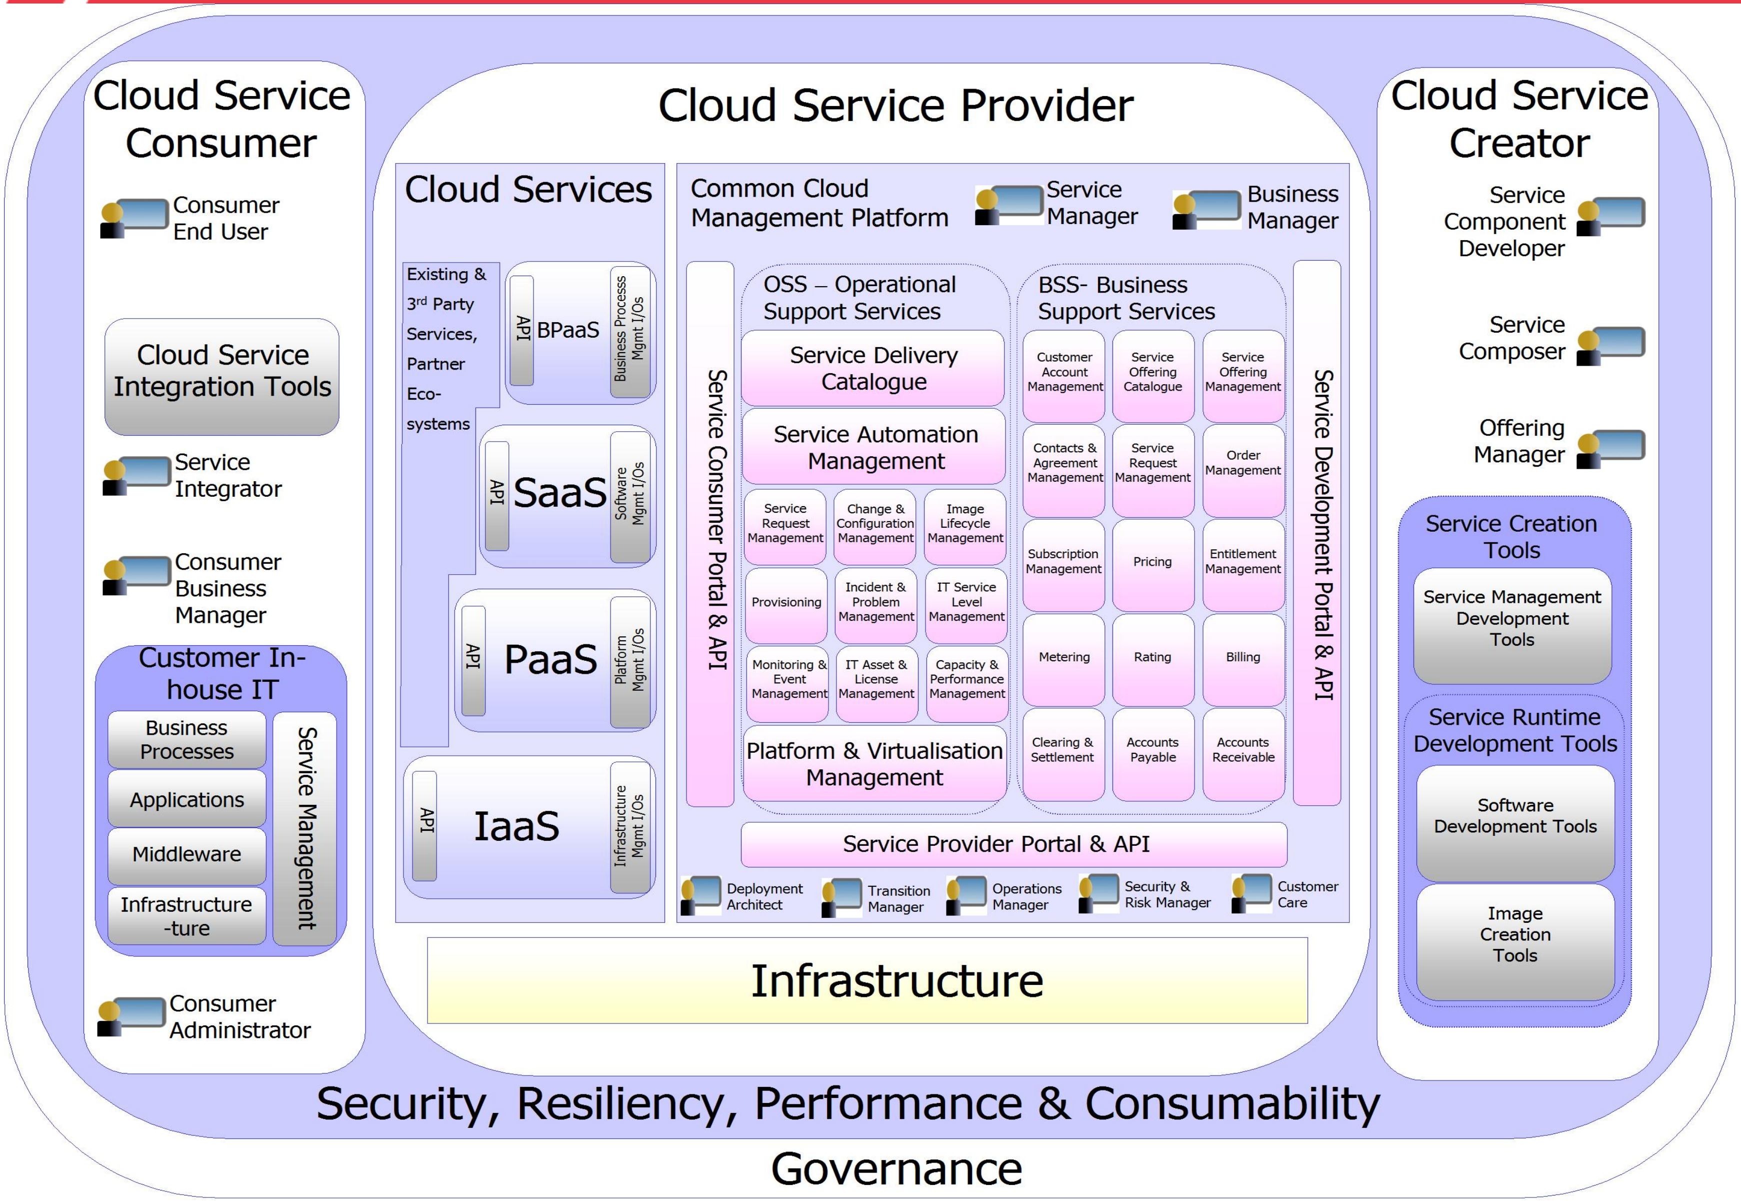

Cloud Service Consumer

Consumer End User

Cloud Service Integration Tools

Service Integrator

Consumer Business Manager

Customer In-house IT

Business Processes

Applications

Middleware

Infrastructure -ture

Service Management

Consumer Administrator

Cloud Service Provider

Cloud Services

Existing & 3rd Party Services, Partner Eco-systems

API BPaaS Business Process Mgmt I/Os

API SaaS Software Mgmt I/Os

API PaaS Platform Mgmt I/Os

API IaaS Infrastructure Mgmt I/Os

Common Cloud Management Platform

Service Manager

Business Manager

Service Consumer Portal & API

OSS – Operational Support Services

Service Delivery Catalogue

Service Automation Management

Service Request Management

Change & Configuration Management

Image Lifecycle Management

Provisioning

Incident & Problem Management

IT Service Level Management

Monitoring & Event Management

IT Asset & License Management

Capacity & Performance Management

Platform & Virtualisation Management

BSS- Business Support Services

Customer Account Management

Service Offering Catalogue

Service Offering Management

Contacts & Agreement Management

Service Request Management

Order Management

Subscription Management

Pricing

Entitlement Management

Metering

Rating

Billing

Clearing & Settlement

Accounts Payable

Accounts Receivable

Service Development Portal & API

Service Provider Portal & API

Deployment Architect

Transition Manager

Operations Manager

Security & Risk Manager

Customer Care

Infrastructure

Cloud Service Creator

Service Component Developer

Service Composer

Offering Manager

Service Creation Tools

Service Management Development Tools

Service Runtime Development Tools

Software Development Tools

Image Creation Tools

Security, Resiliency, Performance & Consumability

Governance

Multi-Year Infrastructure & Delivery Strategy

2009-2011

Public Cloud Adoption

- "Guerilla Cloud Warfare"
- Dev/Test Infrastructure
- Get the company used to them
- Resolve immediate lead time problems

2011-2014

Agile Delivery & Ops

- Move everything to Linux/ Windows
- Agile/lean development
- Automation, configuration management, pervasive virtualization
- Private Cloud for SAP

2012-2015

New Systems Arch

- Fault-Tolerant Distributed DBs & Data Grids
- Event-driven and RESTful integration
- Modular pieces

Public Cloud Adoption

CANADIAN PACIFIC

DRIVING THE DIGITAL RAILWAY »

Scenario: About to hire 200 SAP or Java Consultants

How will you provision for them?

Guerilla Cloud Warfare

- Aka. “How to adopt several hundred desktops & servers in a controlled way with almost no staff”
- **Example Roadblock: Firewalls**
- ***Normal Solution: Open them up.***
  - Discussions, paperwork, pilots, studies, wait 3 months
- ***Guerilla Solution: Reverse SSH Tunnels.***
  Works with TCP, SOCKS, even UDP if you’re crazy enough
- Lesson: Get approval and constraints from the people who matter
  - CIO (who should support your guerilla efforts),
    CISO (who will prepare his team + legal/audit),
    CTO or GM/VP of Architecture (who is supposed to promote new things)
  - Avoid the people who don’t matter, ask forgiveness later

Global Public Cloud Dev/Test Network, late 2010

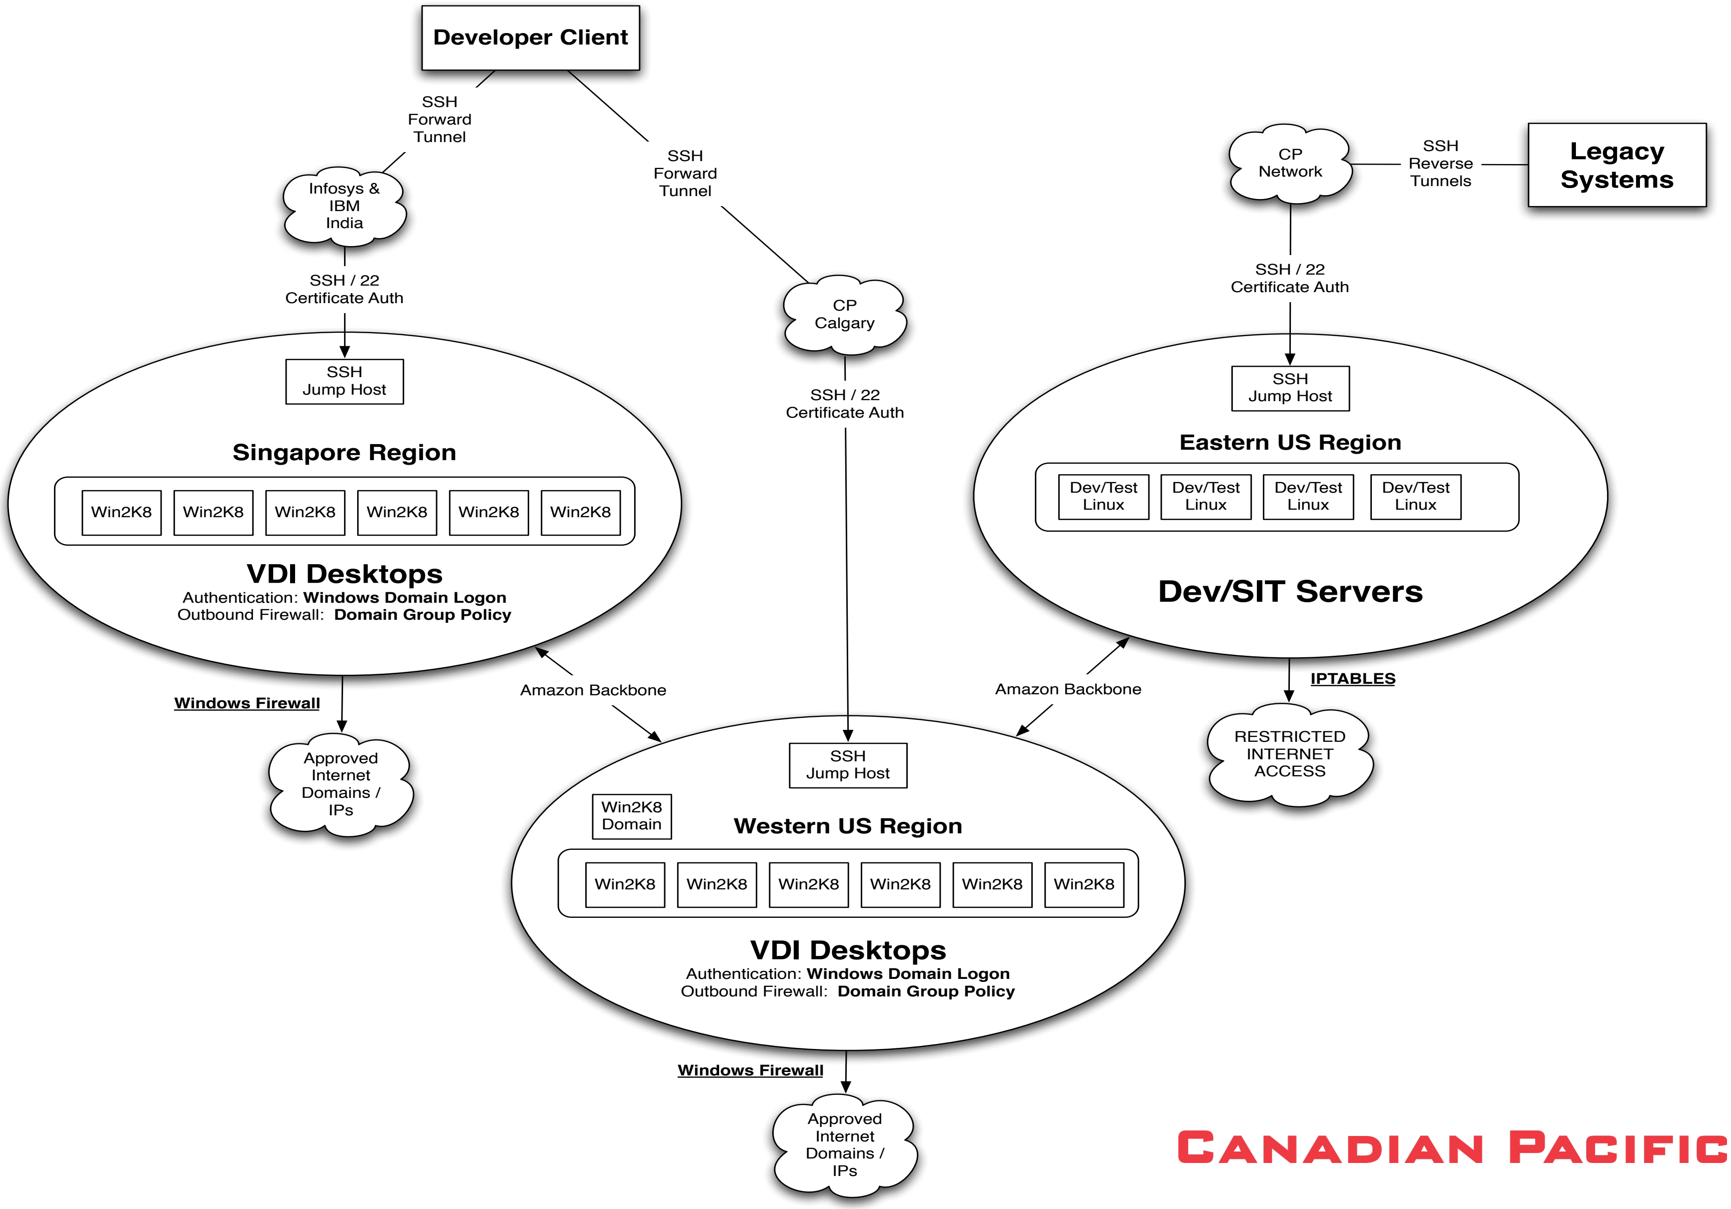

Canadian Pacific

Public Cloud Benefits & Usage Notes

- **Offshore resources get a managed developer workstation**
  - Controlled device admissibility strategy into CP's systems

- **Using Amazon's Internet backbone between regions**
  - More bandwidth, lower latency access to CP's network in Canada
  - Today: Routed via SSH Tunnels
  - Late 2011 / Early 2012: VPN with Overlay Network

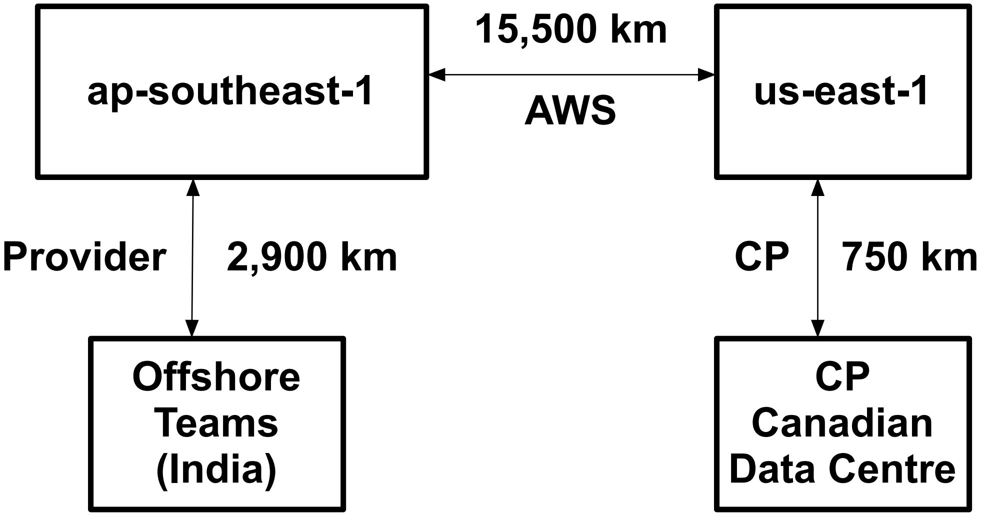

Data Categorization

- **Data Categorization**
  - Handle the legal and regulatory issues associated with data residency
  - Legal desire for physical disks during forensic analysis
  - **Biggest concern: Privacy in the face of a click-through agreement**
  - In short: *Trust your providers* (can't just use "any" cloud provider)

  - Tier 1 Sensitive Data: Harm to Lives (e.g. Hazmat locations)
  - Tier 2 Sensitive Data: Harm to Investors (e.g. financial forecasts)
  - ***Not on public clouds yet***

  - Tier 3 Sensitive Data: Harm to Operations (e.g. Train/car locations)
  - ***On public clouds if in Virtual Private Cloud and encrypted***

  - Tier 4 Sensitive Data: Stale Data and/or Dev/test
  - ***On public clouds***

(Note: These are representative examples, not our actual definitions)

Public Cloud Benefits & Usage Notes

- **Very quick lead times** to deliver working dev/test systems
  - Traditional infrastructure:
    WebSphere, SAP, Business Objects, SQL Server, Exchange, etc.
  - Newer infrastructure: Rails, Haproxy, Nginx, etc.

- **Performance challenges**
  - Most infrastructure clouds do not provide traditionally expected levels of **visibility** in **storage** and **networking**
  - Trend is changing towards ***more visibility & control***
    - E.g. Amazon subnets and routes in VPC
  - **Storage I/O is the major roadblock** to traditional systems
    - E.g. Elastic Block Storage vs. traditional NAS/SAN
    - Latency is not as predictable, node throughput is capped at ~1 Gb, availability is not as predictable

Agile Infrastructure

CANADIAN PACIFIC

DRIVING THE DIGITAL RAILWAY »

Operations: Cultural & Tooling Changes

- **Old Assumptions**
  - "Put your eggs into a small number of baskets, and watch those baskets"

- **New Reality**
  - Partial failure is a regular, normal occurrence;
    no excuse for downtime from any business-level service

- **First Steps to Transformation**
  - Building culture of **collaboration** with IT service delivery
    - Ops offers service engineers as "production service architects"
  - Begin a 5-10 year transition to **"design for failure"** architectures
    - Migration from Mainframe & AIX to Linux (by 2014)
    - In-Memory Data Grids (e.g. WebSphere Extreme Scale)
    - Future: Fault-Tolerant Distributed Databases (e.g. **Riak**)
  - Increasing **visibility** into the operational systems
    - Correlation and drift detection independent of legacy (e.g. **Splunk**)

Enterprise Appliances (Not Really Private Clouds)

- Oracle Exadata
  - Consolidated databases
  - Major OLTP operational data store
  - Major OLAP / data warehouse

- VCE Vblock
  - SAP Landscapes
  - Compute & Midsize DB
  - Exchange

"Wire Once, Walk Away"

Software-Based Automated Configuration

Managed Services that Leverage the Productivity Gains

Private Cloud for Dev/Test

Private Cloud for Production is a Lofty/Questionable Goal

- Thus…

- We're focusing on combining **virtualization** and **appliances** with **automation** & **metrics** to reduce the dev/test cycle

- **CP Application Development & Test Cloud**
  - Vblock + VMware vCloud Director private cloud
    - Pilot Summer 2011, Full Rollout in 2012
  - Linked Clones & Network Fencing for
    - SAP, Legacy, Systems Integration testing
  - Continuing to grow public Cloud Dev/Test Network for new development
    - Continuing with EC2; Piloting vCloud public clouds
  - ITKO LISA for integrated simulation, testing, and validation

Bending the Operational Cost Curve

Projected Monthly Per-Instance Costs (over 3 years)

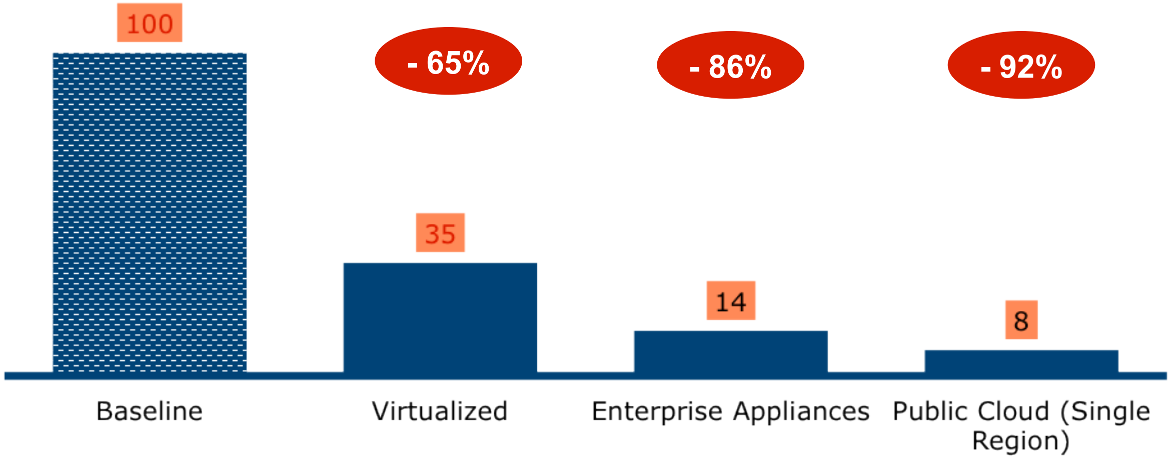

Includes Amortized Capital +
Operating Expense (e.g. Public cloud fees) +
Managed Services

New Systems

CANADIAN PACIFIC
DRIVING THE DIGITAL RAILWAY

The Logic and Constraints of a Railroad

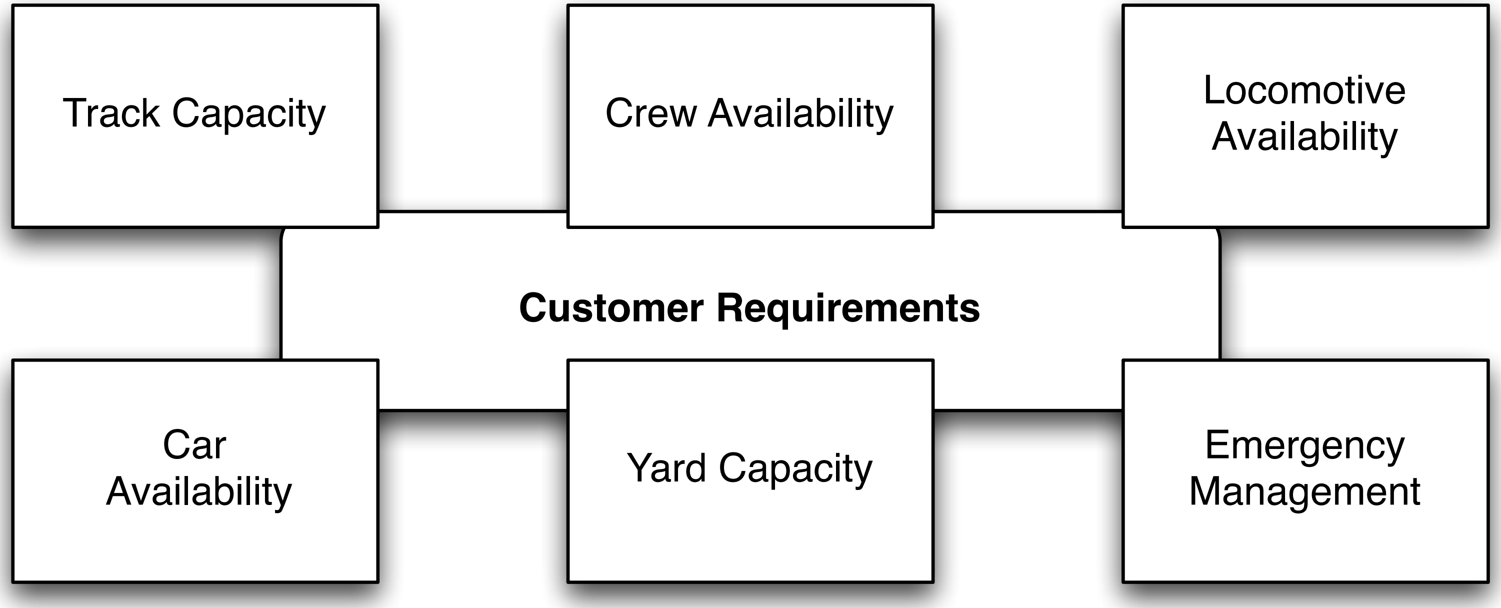

Track Capacity

Crew Availability

Locomotive Availability

Customer Requirements

Car Availability

Yard Capacity

Emergency Management

Basic Railway Systems Architecture (80s)

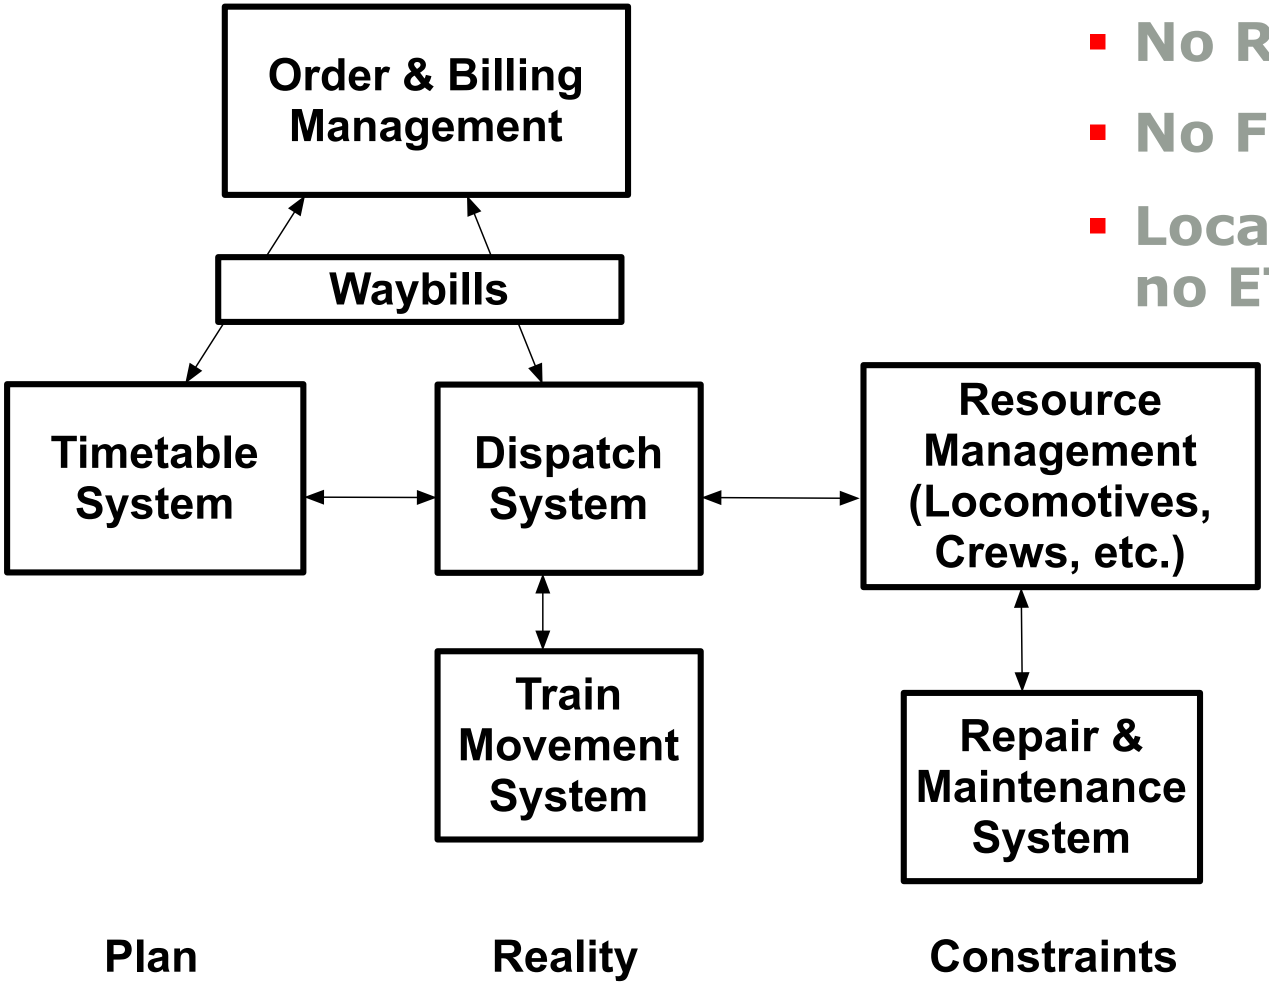

Order & Billing Management

Waybills

Timetable System

Dispatch System

Resource Management (Locomotives, Crews, etc.)

Train Movement System

Repair & Maintenance System

Plan

Reality

Constraints

- No Routing
- No Forecasting
- Location Visibility but no ETAs

Modern Railway System Architecture

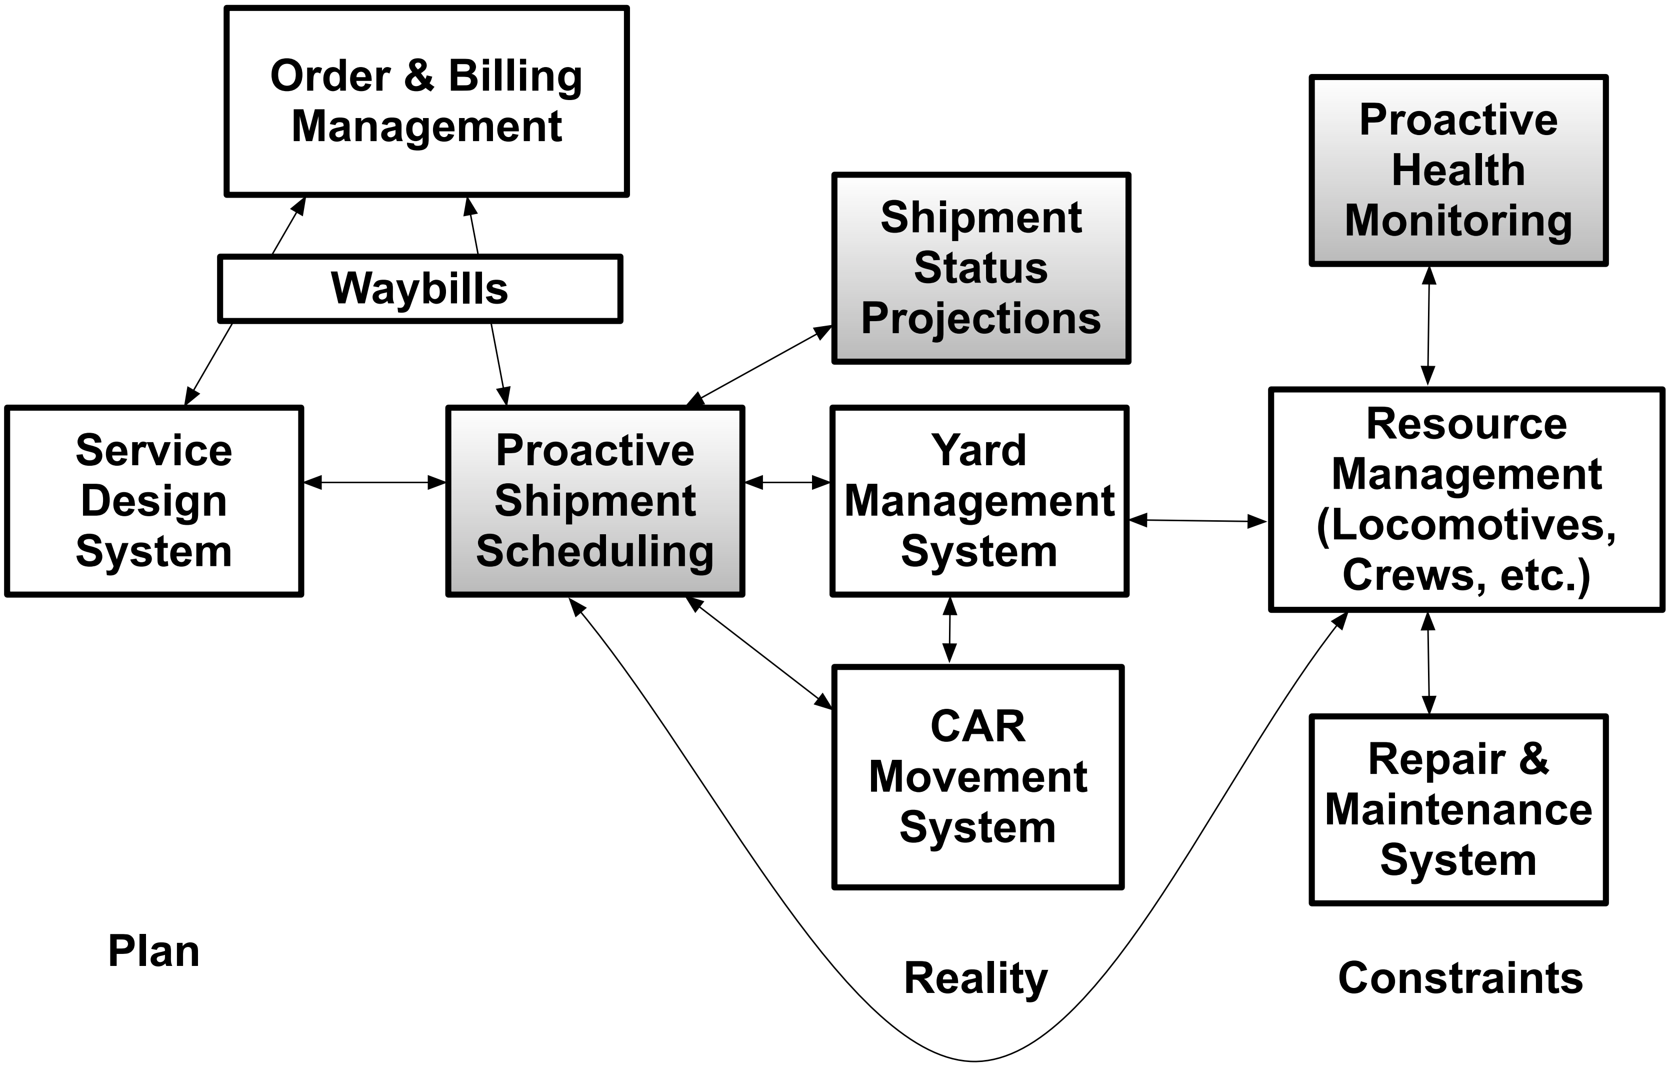
Order & Billing Management
Waybills
Shipment Status Projections
Proactive Health Monitoring
Service Design System
Proactive Shipment Scheduling
Yard Management System
Resource Management (Locomotives, Crews, etc.)
CAR Movement System
Repair & Maintenance System
Plan
Reality
Constraints

Designing a Service, circa 1998-2008

- **Multi-Tier Hybrid Architecture**
  - Some stateless, some stateful computing
  - Session state is replicated
- **Independent** servers / applications
  - Low-level redundancy (RAID, 2x NICs, etc.)
- “Put your eggs into a small number of baskets, and watch those baskets”
- **General assumptions**
  - Failure at the service layer shouldn’t lead to downtime
  - Failure at the data layer may be catastrophic
  - Lots of point-to-point connections
    - ETL, SOAP web services, FTP, etc.

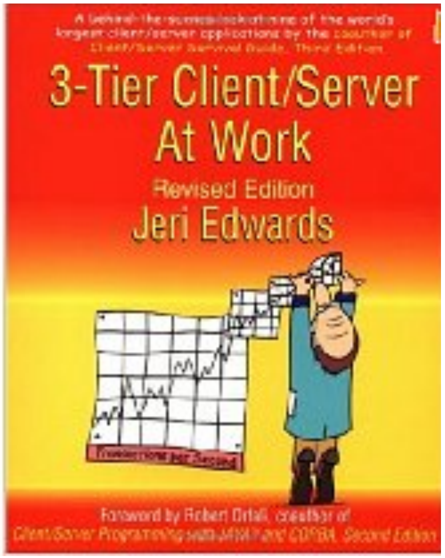
A behind-the-scenes look at some of the world's largest client/server applications by the coauthor of Client/Server Survival Guide, Third Edition
3-Tier Client/Server At Work
Revised Edition
Jeri Edwards
Transactions per Second
Foreword by Robert Orfali, coauthor of Client/Server Programming with JAVA and CORBA, Second Edition

Designing a Service on the Cloud, circa 2008+

- **Autonomous services**
  - Divide system into areas of functional responsibility (tiers irrelevant)
- **Interdependent** servers / applications
  - Software-level redundancy and fault handling
- "Many, many servers breaking big problems down or distributing lots of little problems around"
- **New realities**
  - Partial failure is a regular, normal occurrence; no excuse for downtime from any service
  - Self-describing (RESTful) services for client-device scale
  - Event-driven integration for smaller number of consumers

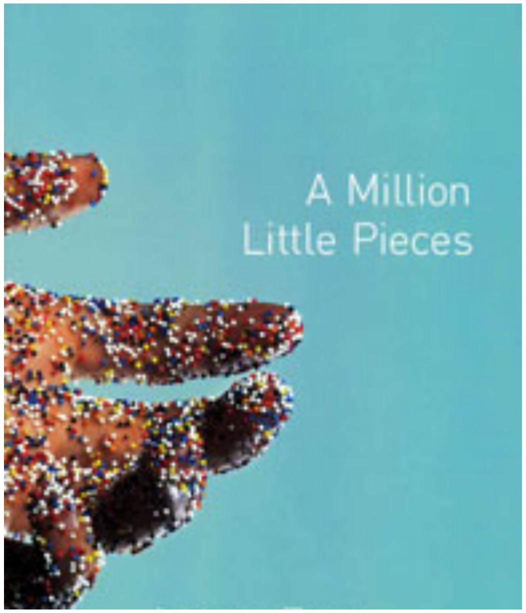

CANADIAN PACIFIC

Current Guidelines for 2012+

Using, where possible: lightweight, simple, inexpensive solutions

1. **High-Performance Event Management** (thousands/sec)
   - Consolidate across multiple proposed event systems
   - Train & Yard Planning, Car Movement, Health Monitoring, PTC
   - Foundation for:
     - Event-Based Integration & predictive real-time analytics
2. **RESTful “Information Resources on Demand”**
   - Self-describing, discoverable, hyperlinked system interfaces & lifecycles
   - No need to directly integrate with databases etc.
   - Foundation for:
     - Business process integration
     - Modern GUIs and Mobile applications
     - Operational BI Mashups
3. **Legacy Endpoint Management**
   - MQ, SOAP Web Services, and Managed File Transfer (EDI)

2012-2015 Systems Design Target (early draft)

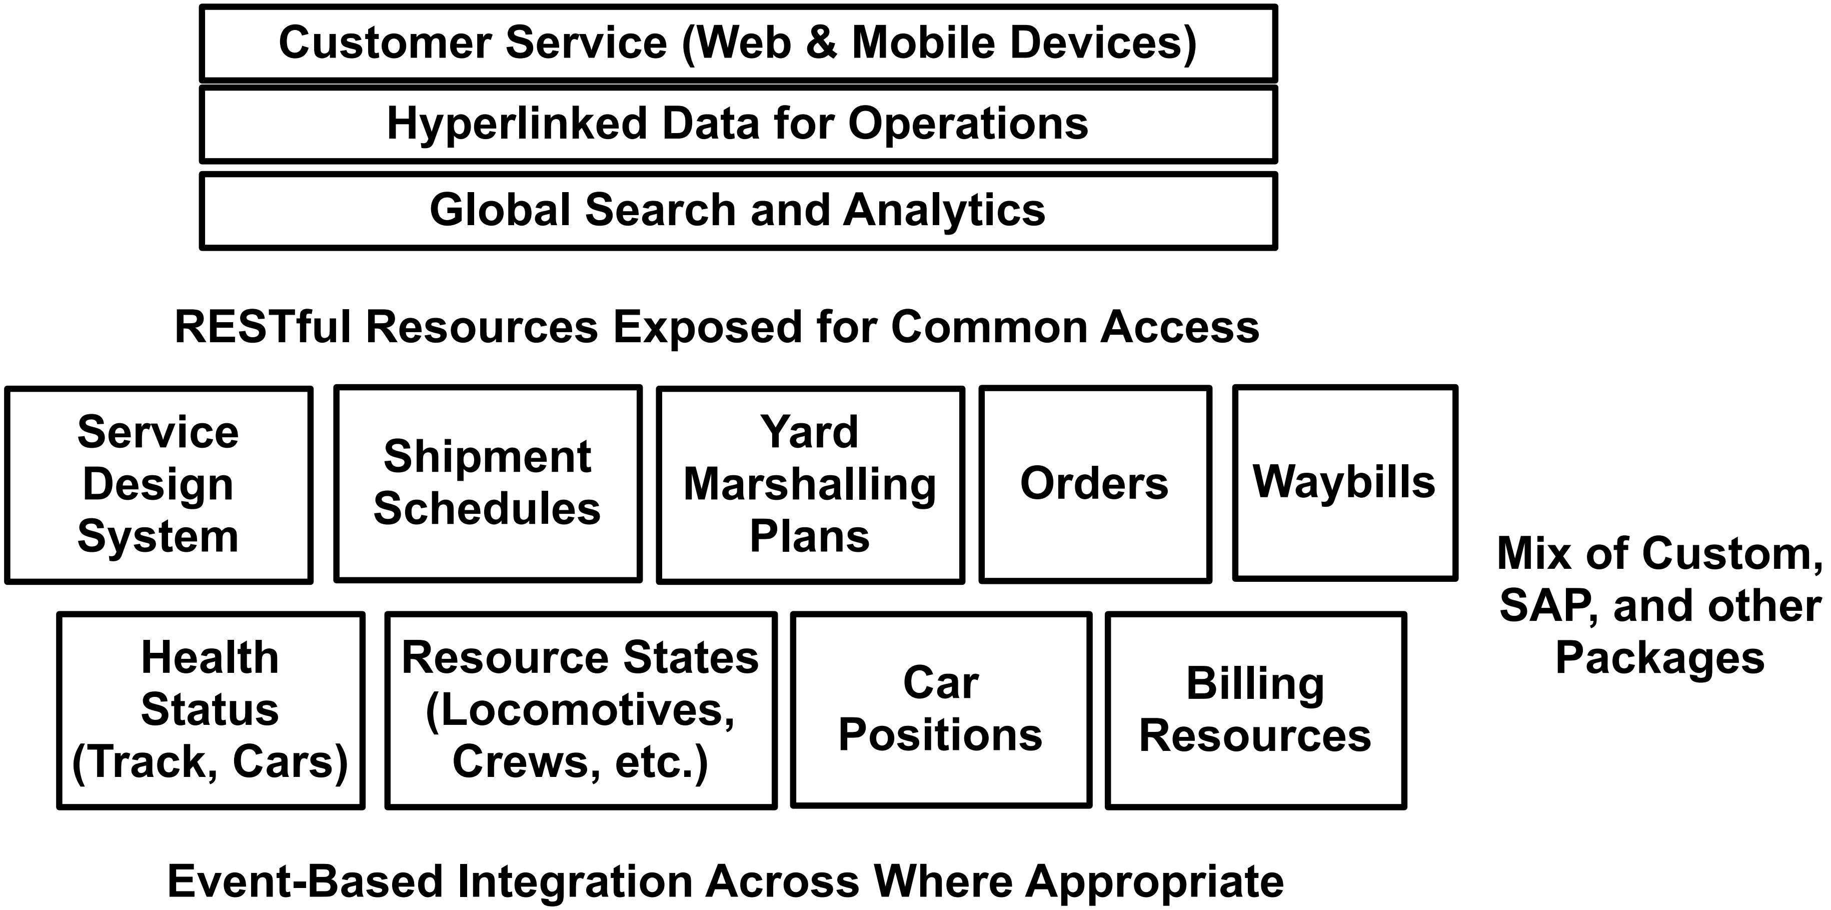

Customer Service (Web & Mobile Devices)

Hyperlinked Data for Operations

Global Search and Analytics

RESTful Resources Exposed for Common Access

Service Design System

Shipment Schedules

Yard Marshalling Plans

Orders

Waybills

Health Status (Track, Cars)

Resource States (Locomotives, Crews, etc.)

Car Positions

Billing Resources

Mix of Custom, SAP, and other Packages

Event-Based Integration Across Where Appropriate

Summary:
Multi-Year Infrastructure & Delivery Strategy

2009-2011

Public Cloud Adoption

- “Guerilla Cloud Warfare”
- Dev/Test Infrastructure
- Get the company used to them
- Resolve immediate lead time problems

2011-2014

Unified Infrastructure

- Move everything to Linux/ Windows
- Agile/lean development
- Automation, configuration management, pervasive virtualization
- Private Cloud for SAP

2012-2015

New Systems Arch

- Fault-Tolerant Distributed DBs & Data Grids
- Event-driven and RESTful integration
- Modular pieces

Contacts & Thanks

Stuart Charlton

Director – Infrastructure & Operations

Information Technology

Stuart_Charlton@cpr.ca

Canadian Pacific
Suite 500, 401 – 9th Avenue SW
Calgary Alberta
Canada T2P 4Z4
www.cpr.ca

With thanks to….

CP architecture: Gary Stedman, Dragan Sajic, Vincent Blue, Tim Riley

CP operations: Bob Nash, Jack Vanos, Michael Turcotte, Ron Legere, Stan Singer

CP IT risk management & security: Kevin Pasveer

CP application delivery: Shawn Adams, Michael Wiens, Steve Hester

CP CIO: Heather Campbell