



ONE COMMON PLATFORM.

 **DealerTrack**[®]

**NEW
RULES**

Linux as a Data Integration Platform

**BETTER
TOOLS**

**Evan Bauer, VP Technology Architecture
DealerTrack, Inc.**

 **DealerTrack**[®]
NEW RULES **BETTER TOOLS**



- DealerTrack (NASDAQ:TRAK) is the leading software-as-a-service and data provider to the automotive industry
- Origins as the spinout of Chase's auto finance portal
- In our 10-year history, DealerTrack has made 17 acquisitions and developed 4 major new products internally
- Integration of our products and the ability to create new resources from the databases they build and create is the company's #1 strategic project
- Stabilization of the runtime environment for availability and increasing scalability to handle rapid growth in transaction volumes is part of the price of success in the market

What are we integrating?



- A combination of interactive web, web services, and 5250 telnet delivered SaaS products for the retail automotive supply chain
- Customers include:
 - Dealers
 - Lenders
 - Manufacturers
 - Insurers
 - Parts and Service Suppliers
 - State Governments
- Implemented with:
 - IIS / .NET / MSMQ
 - Java / Weblogic / J2EE
 - Apache / perl
 - RPG2 / LookSoftware
 - Oracle
 - MS SQL Server
 - DB2
 - MySQL

The screenshot shows a web browser displaying the DealerTrack website. The browser's address bar shows the URL: <https://www.dealtrack.com/public/login.fcc?TYPE>. The website header includes the DealerTrack logo and navigation links: HOME | CONTACT US | USA | CANADA. Below the header is a navigation menu with links for Solutions, Financing Sources, Partners, Support & Training, and About Us. The main content area features a large advertisement for DealerTrack DMS. The ad includes a padlock icon and the text: "Ask us about our 48 month DMS price-lock guarantee." and "Lock in long-term savings" and "No long-term contract". Below this, it states: "Save 50% by switching to the #1 web-based DMS, and lock in a low fixed monthly payment for 48 months.*" and "Learn how to lock in savings." There is also a "Limited time offer." and "DealerTrack Performance Suite" section with buttons for DMS, INVENTORY, SALES AND F&I, and COMPLIANCE. On the right side of the page, there is a "Login" section with fields for Login ID and Password, a "Log In" button, and a "Forgot Password?" link. Below the login section, there is a "Attention: Toyota Dealers" section with text: "DealerTrack DMS now features full integration with Toyota Dealer Daily." and a "Learn more now." button. At the bottom of the page, there is a "Find Lenders Serving Your Area" section with a "Enter Your Zip Code:" field and a "Lookup" button. The footer of the website includes links for Terms and Conditions, Privacy Policy, Site Requirements, Investor Relations, Careers, and Contact Us, as well as social media icons for Blogs, Twitter, and YouTube. There is also a "SECURITY MATTERS" section with a link to "Learn how to keep your dealer information secure" and a "Verizon CyberTrust Security" logo.



Constraints:

- Business is sensitive to both availability and cost
- Wholesale application re-hosting or conversion is much too costly and would provide marginal user benefit
- Subject matter expert developers are our scarcest resources and are often deeply wedded to specific development environments
- Integration of data is a prerequisite to integrating applications
- Thousands of external data integration points with 3rd parties

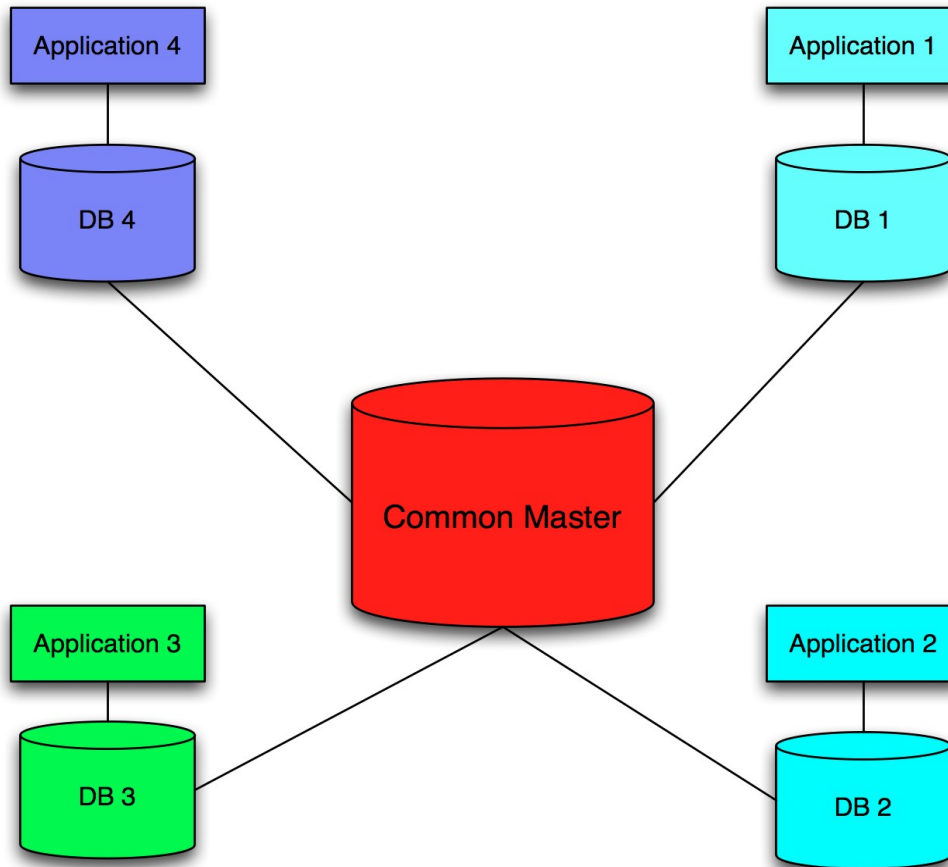
Approach:

- Definition of common data entities that are shared across platforms
- Minimum modification of applications to create and consume shared data
- Data replication between databases in a hub and spoke topology



- Create a common data platform for all current and future DealerTrack Products and Services
- Design and implement a multi-platform software/Database framework upon which DT 2.0 will be built
- Create and maintain database tables that will store common data and make it easily available to all products regardless of platforms
- Implement replication software so that common data will be replicated to all the platforms and will be available locally to all products
- Common Entities to be released in 2011 include: Dealer Groups, Dealers, Users, Customers, Vehicles, and Partners

Fusion Replication Topology





- Change Data Capture selected as the replication mechanism
- Data replication becomes a service provided by the DBA team, rather than the programmatic responsibility of the application development teams
- Replications defined as the relationships between tables, with both copies of the Common Master available in application DBs (full or in subset) as well as replication with transformation to modified legacy tables
- Approach simplifies and accelerates delivery of data integration
- Common Master – the hub – becomes a hotspot. It receives the sum of all updates and needs to be available as the projection of all application SLAs



SLA and the Clock

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

<u>Repair Time</u>		<u>Availability</u>
8 hours	$\frac{1000}{1000 + 8}$	= 99.2 %
½ hour	$\frac{1000}{1000 + 0.5}$	= 99.95 %
20 seconds	$\frac{1000}{1000 + 20/3600}$	= 99.999%

Assumptions:

1. Mean Time to Repair is Controllable with tools and architecture
2. Mean Time Between Failure shown at 1,000 hours – once every 40 days

Availability is sensitive to MTTR – even 3-nines is hard to reach without having high availability during maintenance windows. Rolling upgrades and ≥ 3 -way live-live clusters are essential for stateful subsystems



- Original architecture for production OLTP systems had stateless web farms in front of Solaris SPARC Oracle database servers on Veritas clusters
- Lower cost Linux clusters on commodity x86 servers, initially implemented for less critical or lower volume systems, have had substantially better availability and performance
- The fault-rate associated with system software root causes has been lower on Linux-based clusters than on either Solaris or Windows
- Measured MTTR is lower on our Linux clusters than on any other platform in production
- Device driver support, a critical cause of outages and a substantial engineering and testing burden on our Windows IIS and MSMQ servers, has never been an issue on Linux database clusters running on comparable hardware from the same manufacturer



- As desirable as a fully open source solution would be, support for both Windows and iSeries (OS/400) legacy requires proprietary CDC and RDBMS technology
- Without at least 3 active cluster nodes, you don't have HA during maintenance and upgrades
- You can't use DR as HA when 99.99% uptime is your target
- For all vendors of alternatives for the necessary software, Linux has become the common, tier-1, supported platform
- Key issues become file system and support across SAN alternatives and the interaction between CDC, backup, and DR storage replication technologies
- Testing and stabilization become as critical as design and development



Once the platform decision was made, additional cost-effective Linux-based solutions presented themselves for related systems:

- Replication of transactional data to operational data stores for to support ad hoc reporting. Linux versions of both reporting and database software provide a 4x price performance advantage over the same proprietary software running on iSeries
- Cost-effective av-scanned, remote replicated, encrypted storage of sensitive documents
- Host platform for utility services (e.g. network management, log concentration and analysis) supporting tools for the diverse set of primarily proprietary platforms

Key Future Considerations



Linux as a database and infrastructure service delivery platform has been adopted within DealerTrack because of the strengths that come from an open platform – despite fears about “open”.

In turn, Linux is helping to open DealerTrack's platform and options going forward.

- RAS – there is nothing worse than being offline
- RDBMS cluster support facilities for live-live clusters (RAC or PureScale vs HA or HA-DR)
- “Cookbook” filesystem selection and implementation for database clusters
- LSB remains critical as database and proprietary ISV products maintain distro affinities
- Virtualization – kvm becoming increasingly attractive, proprietary offering feature rich but problematic in production. Management tools welcome