

*Exceptional service in the national interest*



# IDC Reengineering Phase 2

## Prototyping Status

Ryan Prescott

22 June 2015

SAND Number:



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Agenda

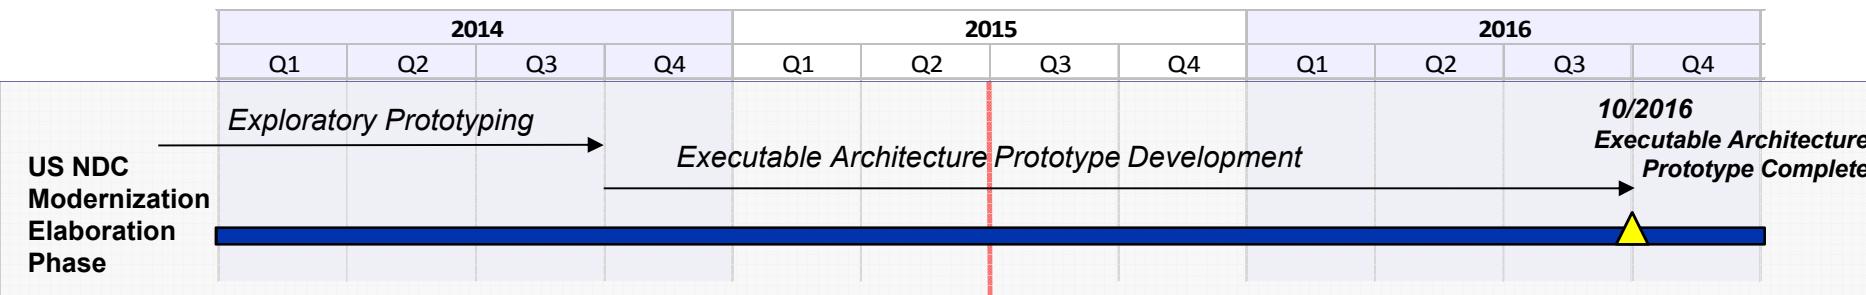
- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

# Prototyping Overview (1 of 2)

- The US NDC Modernization project plan includes a software prototyping component supporting definition of the system architecture
- Prototyping is intended to facilitate:
  - Definition of high-level design patterns
  - Demonstration of key architecture concepts & features
  - Selection of representative technologies
    - System platform
    - Software languages
    - Third-party software

# Timeline

- The project plan includes two prototyping phases:
  - Exploratory Prototyping, FY2014
    - Focused on technology evaluation
    - Software language selection and third-party software evaluations for key software mechanisms
    - See backup slides for technology evaluation summary
  - Executable Architecture Prototyping, FY2015-FY2016
    - Focused on demonstration of key system features and mechanisms
    - Following a SCRUM process with 3-week sprints



# EXECUTABLE ARCHITECTURE PROTOTYPE

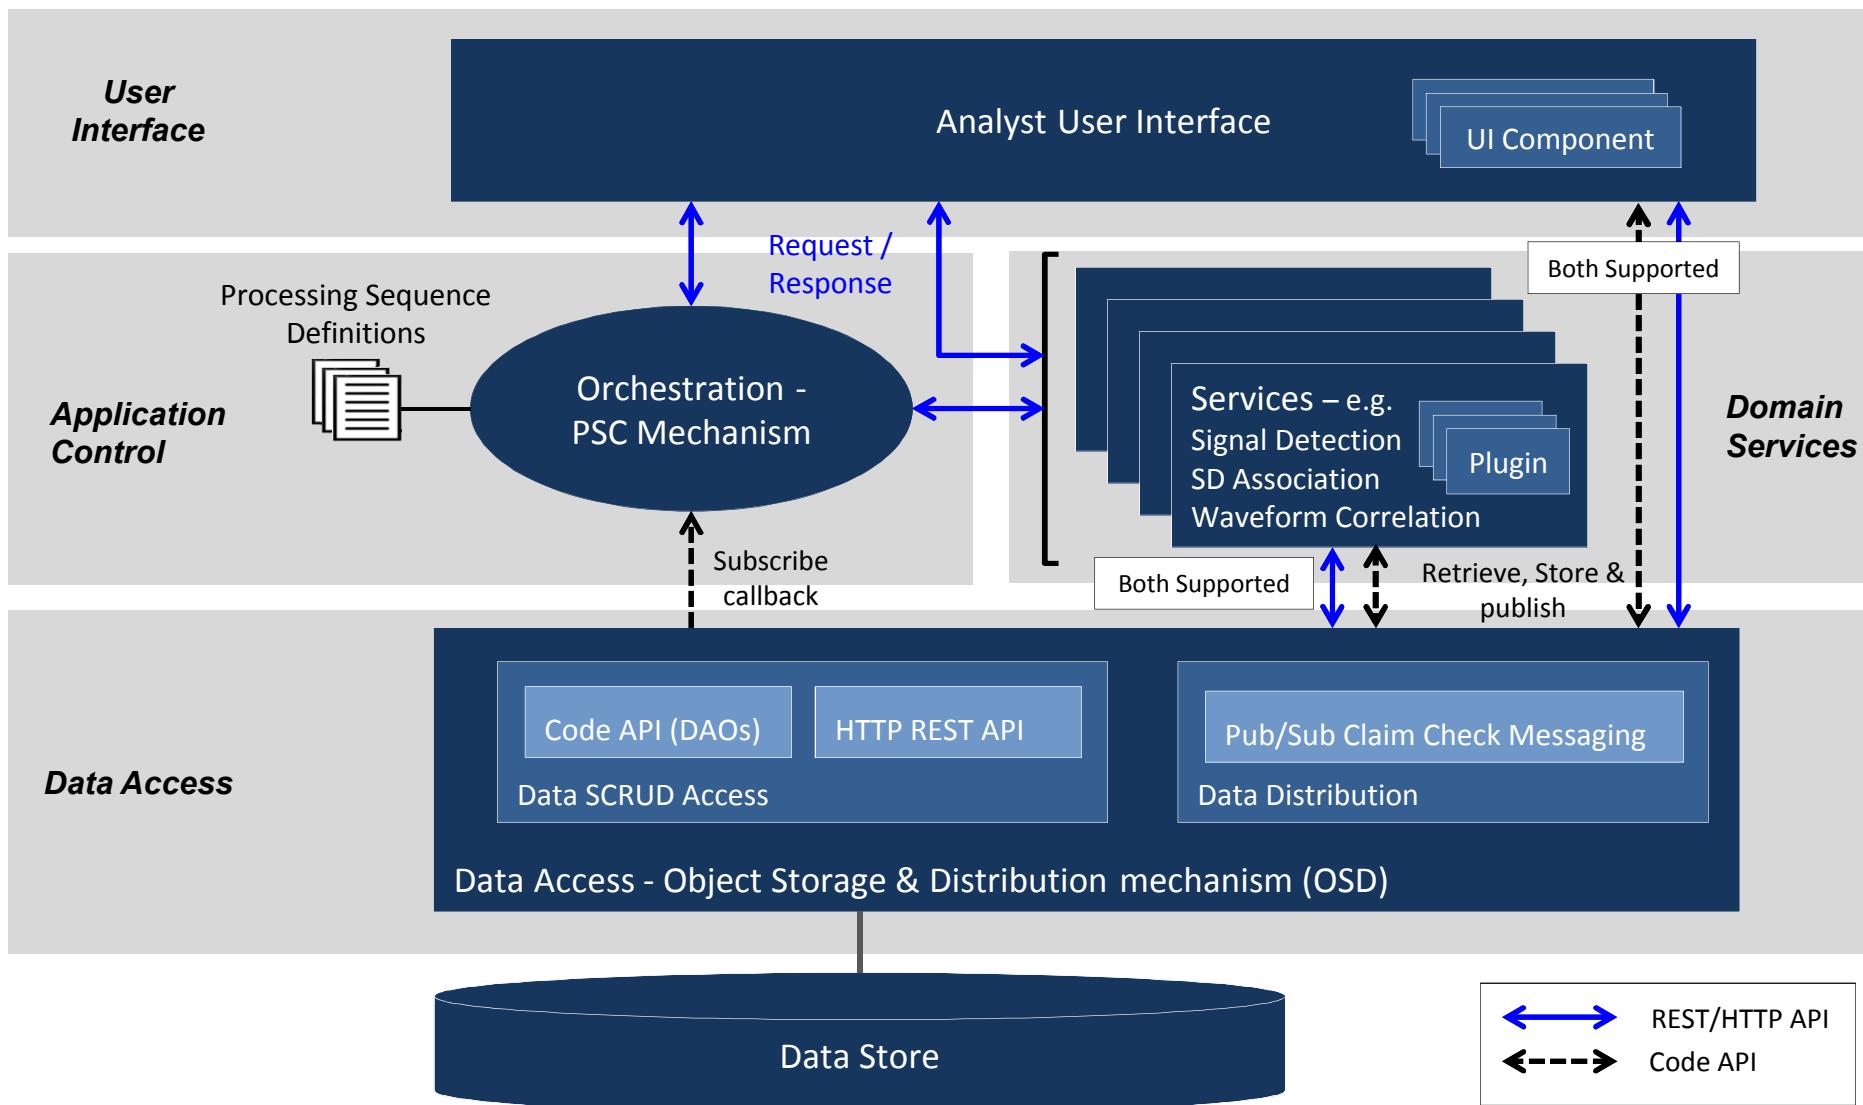
# Agenda

- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

# Executable Architecture Prototype – Definition

- Elaboration-phase activity to implement a portion of the system architecture as defined in the Architecture Document and Analysis Model
- Determine if it is feasible to implement the architecture
  - Must satisfy system requirements
  - Feedback loop to update architecture when needed
- Executable Architecture is a prototype
  - The intent is to validate key features & mechanisms of the architecture rather than to develop an early version of the system

# Executable Architecture Prototype – Conceptual Overview



# Executable Architecture Prototype –

## Assumptions (1 of 3: platform and languages)

- **Platform:**
  - Infrastructure:
    - Distributed deployment: bare metal, private cloud (e.g. OpenStack), possibly Cloud
  - OS: Centos 7+ (current) / RHEL 7+ (future)
- **Storage architecture:** RDBMS (Oracle, Postgres, etc.)
- **Software Development Languages:**
  - User Interface: Java &/or JavaScript (open trade re: desktop vs. browser-based UI)
  - Application Control & Orchestration: Java
  - Domain Services:
    - Multiple supported (C++, Java, Python)
    - Only Java demonstrated
  - Data Access:
    - Code API: Java (C++ TBD)
    - REST API: Multiple supported (C++, Python, Java, etc.), Java included in prototype

# Executable Architecture Prototype –

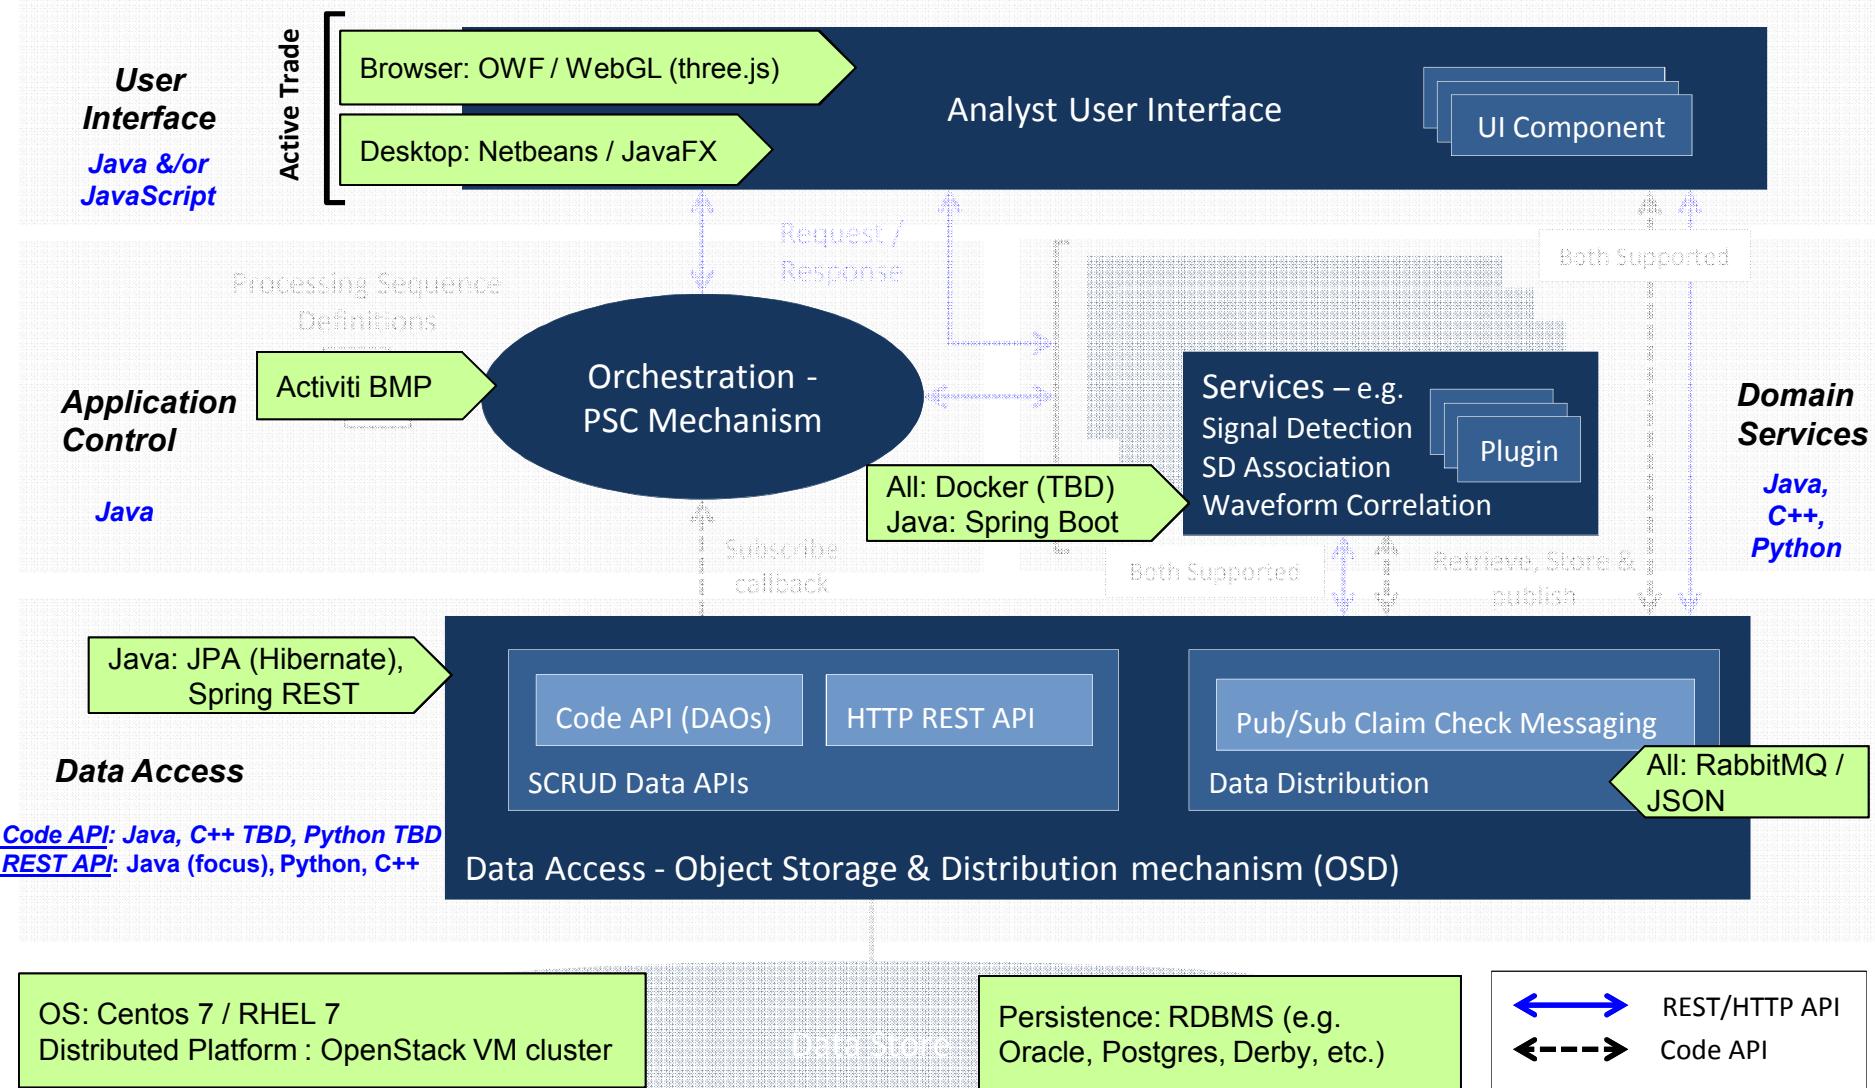


## Assumptions (2 of 3: third-party software)

- **User Interface:**
  - Browser-Based
    - Framework: Ozone Widget Framework
    - Waveform displays: OpenGL (three.js)
    - Map displays: TBD (Worldwind, Google Maps)
  - Desktop
    - Framework: Netbeans RCP
    - Waveform displays: JavaFX charts, OpenGL (TBD)
    - Map displays: TBD (Worldwind, Google Maps)
- **Processing Sequence Controller mechanism (PSC):** Activiti BPMN Engine
- **Domain Service Encapsulation:** docker (TBD - see backup for docker overview)
- **Inter-process Communications:**
  - Data Distribution (pub/sub): RabbitMQ / JSON
  - Service Invocation (request/response): HTTP(S) / JSON (REST)
    - Java: Spring REST
- **Data Access (SCRUD):**
  - Code API - Java: JPA (Hibernate), C++ (support TBD)
  - REST API: Multiple clients supported Spring REST included in prototype
- **Data Access (Distribution): RabbitMQ / JSON**

# Executable Architecture Prototype –

## Assumptions (3 of 3: Technology Summary)



# Executable Architecture Prototype – Key Features

- Implemented features demonstrate fundamental concepts
  - PSC: processing sequence definition, execution, control
  - OSD/COI: data models, persistence, data distribution, data provenance
  - UI: modern frameworks, extensible, undo/redo, OSD/PSC integration
- Features demonstrate high risk architectural aspects and non-functional (“-ility”) SRDs
  - Usability
    - Low-latency, high-performance analysis user interfaces
    - Show undo/redo, data synchronization, user customizable displays
  - Data Provenance
    - Tracking and preserving data availability, processing parameters, processing histories, and the users/processes who worked on all results
  - Configurability
    - Processing parameters configurable by station, phase, etc.
    - Processing sequences initiated based on configurable criteria
  - Maintainability / Extensibility
    - Creating data abstraction layers and pluggable algorithm implementation patterns
  - Deployment
    - A variety of deployments are required. Will demonstrate a data center deployment supporting local and remote interactive analysis.
- **Approach: Demonstrate requirements through the implementation of select user scenarios**

# User Scenario (1/3) – Interactive Analysis

- Analyst selects data (time interval or selected event)
- Views waveforms, signal detections, and events
- Interacts with waveforms
  - Scroll, pan, zoom, scale, and filter
- Works with signal detections and events
  - Create signal detections, create events, modify certain signal detection/event parameters
- Undo / redo certain operations
- Marks data analysis complete
  - Mark processing stage complete
- Receives notifications about other Analyst or System activity
  - Option to update display to show changes
- Views Data provenance for events & signal detections
- Alternate scripting interface to access waveform, signal detection, event hypotheses, etc.

# User Scenario (2/3) – OSD support

- Implement database access abstraction
  - SCRUD Java Code APIs implemented in Java DAOs
  - Partial entity classes developed for waveforms, signal detections, events, processing sequences, and processing stages
  - REST-ful HTTP / JSON SCRUD APIs partially implemented to evaluate performance
- Implement data distribution (pub/sub) both with serialized Entity classes & using the *claim check* pattern
  - Publications may trigger
    - Processing sequence execution
    - Notifications to other Analysts
- Implement data model and persistence updates in order to record and display provenance and event history

# User Scenario (3/3) – PSC support

- Define and execute mock processing sequences for automated processing of signal detections and event hypotheses
  - Processing Sequences predefined and persisted in textual format (*BPMN 2.0 XML standard*)
  - OSD loads Processing Sequences into Entity classes
  - Use stubbed domain services accessing stubbed Plugins
  - Processing sequences executed based on pub/sub data distribution and direct REST-based invocation
    - Implement interfaces based on Entity classes
    - Geophysics algorithm implementations possible, not required
- Demonstrate Processing Sequences triggered by OSD callbacks, e.g.
  - Analyst modifying Signal Detection attribute
  - Analyst modifying an Event attribute
  - Analyst marking processing stage complete

# CURRENT PROTOTYPE STATUS

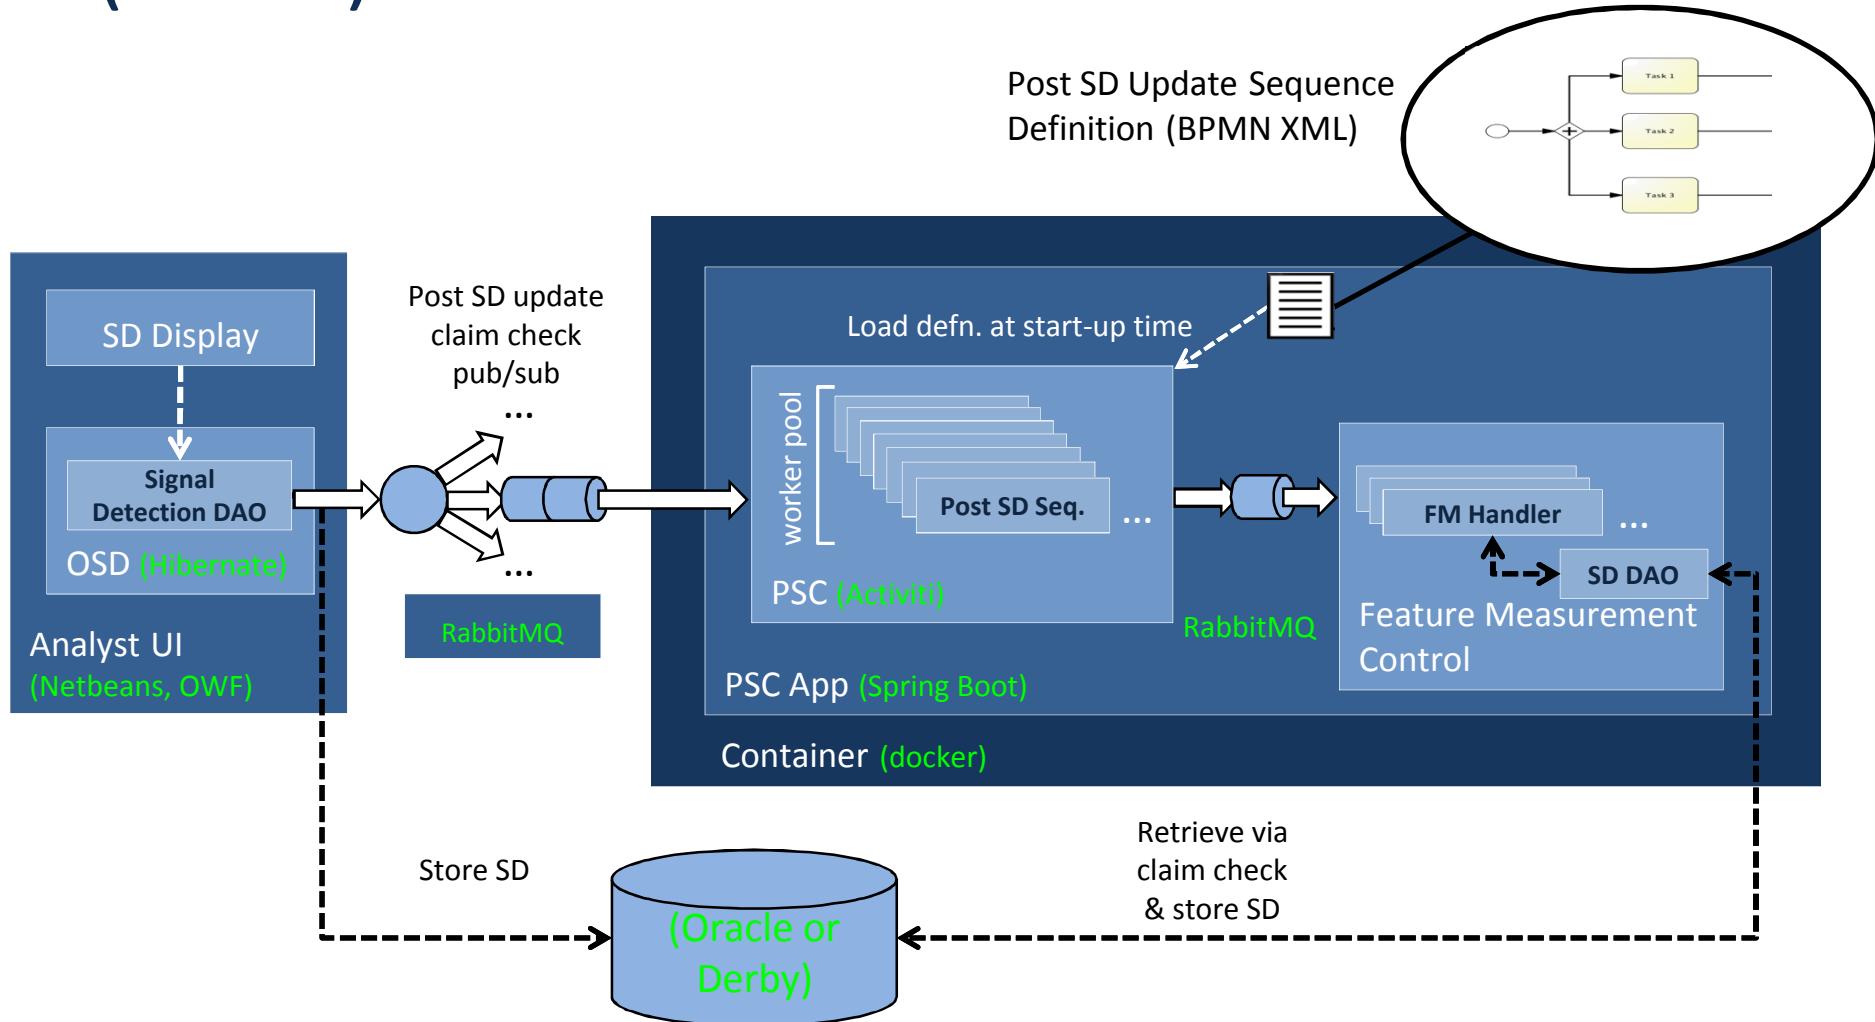
# Agenda

- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

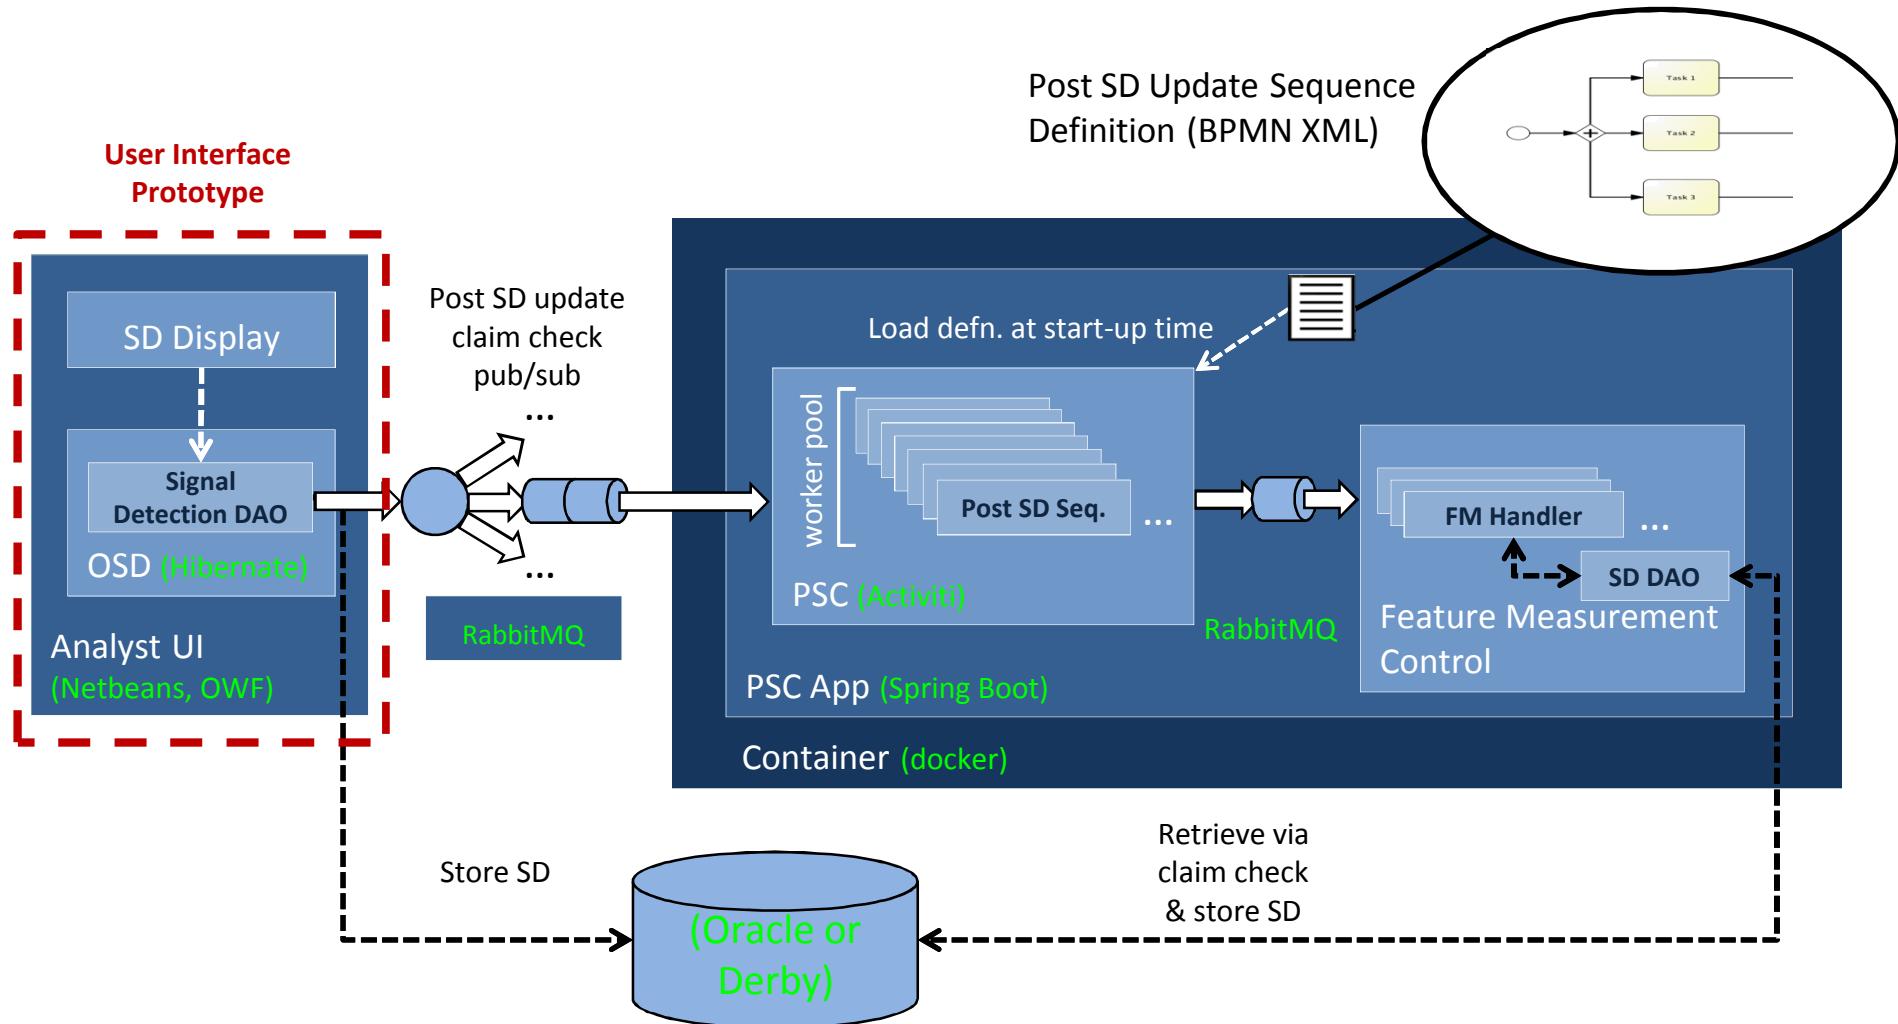
# Current Prototype Status - Overview (1 of 2)

- Focus: Signal detection automated processing and analysis
- User Interface: Waveform display and signal detection table
  - Scroll, pan, zoom and scale waveforms
  - Create & re-time phase-labeled signal detections, storing and publishing new/updated signal detection entities via the OSD API
  - Competing desktop and browser-based implementations (Netbeans and OWF, respectively)
- Application Control: Automated execution of mock signal detection processing sequences via the PSC mechanism
  - Selected COTS BPMN engine (Activiti)
  - Developed an initial PSC prototype
- Object Storage & Distribution: Storage/retrieval and distribution of signal detection & waveform entities
  - Load waveform data into the display via the OSD
  - Store and retrieve signal detection entities via the OSD
  - Notify subscribers via pub/sub (RabbitMQ) whenever signal detections are created or modified

# Current Prototype Status - Overview (2 of 2)



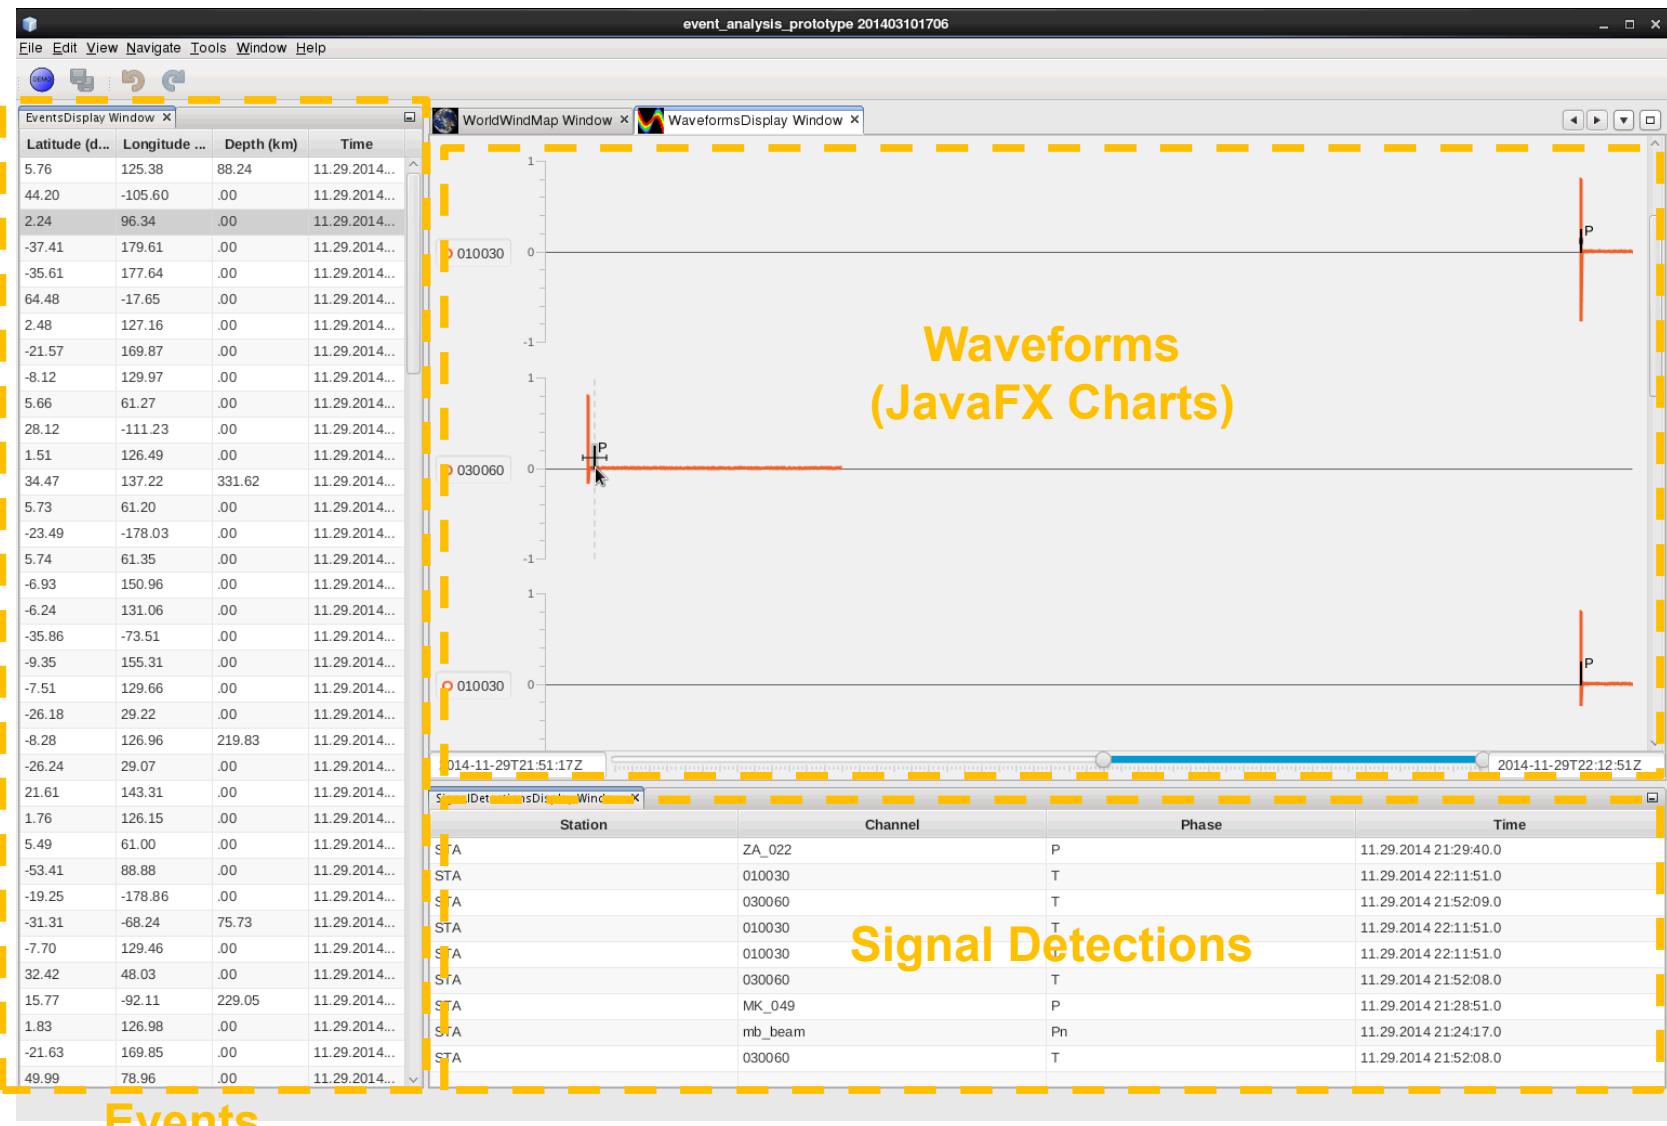
# User Interface Prototyping Status (1 of 2)



# User Interface Prototyping Status (2 of 2)

- Developed initial Netbeans/JavaFX signal detection displays
  - Implemented a waveform plotting display
  - Implemented a map display
  - Implemented event and signal detection table displays
- Developed initial OWF/WebGL signal detection displays
  - Implemented a waveform plotting display
  - Implemented a map display

# Netbeans Display Prototype – Signal Detection (1 of 2)



# Netbeans Display Prototype – Signal Detection (2 of 2)

event\_analysis\_prototype 201403101706

File Edit View Navigate Tools Window Help

EventsDisplay Window X

Latitude (d... <th>Longitude ...<th>Depth (km)<th>Time</th></th></th>	Longitude ... <th>Depth (km)<th>Time</th></th>	Depth (km) <th>Time</th>	Time
5.76	125.38	88.24	11.29.2014...
44.20	-105.60	.00	11.29.2014...
2.24	96.34	.00	11.29.2014...
-37.41	179.61	.00	11.29.2014...
-35.61	177.64	.00	11.29.2014...
64.48	-17.65	.00	11.29.2014...
2.48	127.16	.00	11.29.2014...
-21.57	169.87	.00	11.29.2014...
-8.12	129.97	.00	11.29.2014...
5.66	61.27	.00	11.29.2014...
28.12	-111.23	.00	11.29.2014...
1.51	126.49	.00	11.29.2014...
34.47	137.22	331.62	11.29.2014...
5.73	61.20	.00	11.29.2014...
-23.49	-178.03	.00	11.29.2014...
5.74	61.35	.00	11.29.2014...
-6.93	150.96	.00	11.29.2014...
-6.24	131.06	.00	11.29.2014...
-35.86	-73.51	.00	11.29.2014...
-9.35	155.31	.00	11.29.2014...
-7.51	129.66	.00	11.29.2014...
-26.18	29.22	.00	11.29.2014...
-8.28	126.96	219.83	11.29.2014...
-26.24	29.07	.00	11.29.2014...
21.61	143.31	.00	11.29.2014...
1.76	126.15	.00	11.29.2014...
5.49	61.00	.00	11.29.2014...
-53.41	88.88	.00	11.29.2014...
-19.25	-178.86	.00	11.29.2014...
-31.31	-68.24	75.73	11.29.2014...
-7.70	129.46	.00	11.29.2014...
32.42	48.03	.00	11.29.2014...
15.77	-92.11	229.05	11.29.2014...
1.83	126.98	.00	11.29.2014...
-21.63	169.85	.00	11.29.2014...
49.99	78.96	.00	11.29.2014...

WorldWindMap Window X

WaveformsDisplay Window X

SignalDetectionsDisplay Window X

Each tab is a “top component” that can be docked, undocked, closed, reopened, resized, etc.

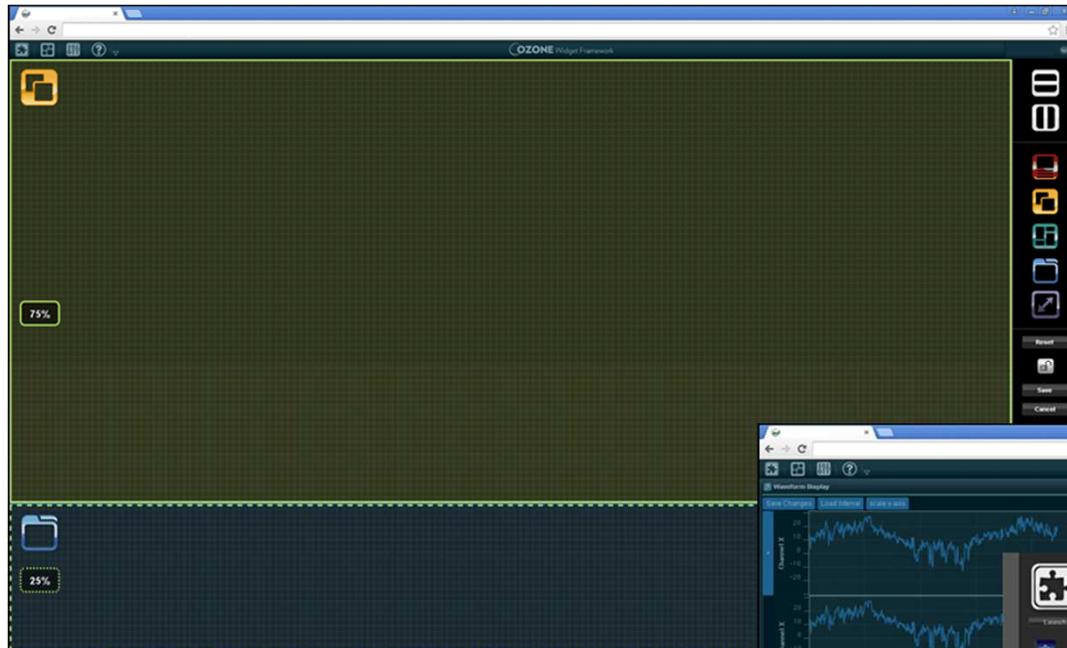
NetBeans remembers user layout settings between sessions.

2014-11-29T21:51:17Z

2014-11-29T22:12:51Z

Station	Channel	Phase	Time
STA	ZA_022	P	11.29.2014 21:29:40.0
STA	010030	T	11.29.2014 22:11:51.0
STA	030060	T	11.29.2014 21:52:09.0
STA	010030	T	11.29.2014 22:11:51.0
STA	010030	T	11.29.2014 22:11:51.0
STA	030060	T	11.29.2014 21:52:08.0
STA	MK_049	P	11.29.2014 21:28:51.0
STA	mb_beam	Pn	11.29.2014 21:24:17.0
STA	030060	T	11.29.2014 21:52:08.0

# OWF Display Prototype – Default and Custom Dashboards (2 of 3)

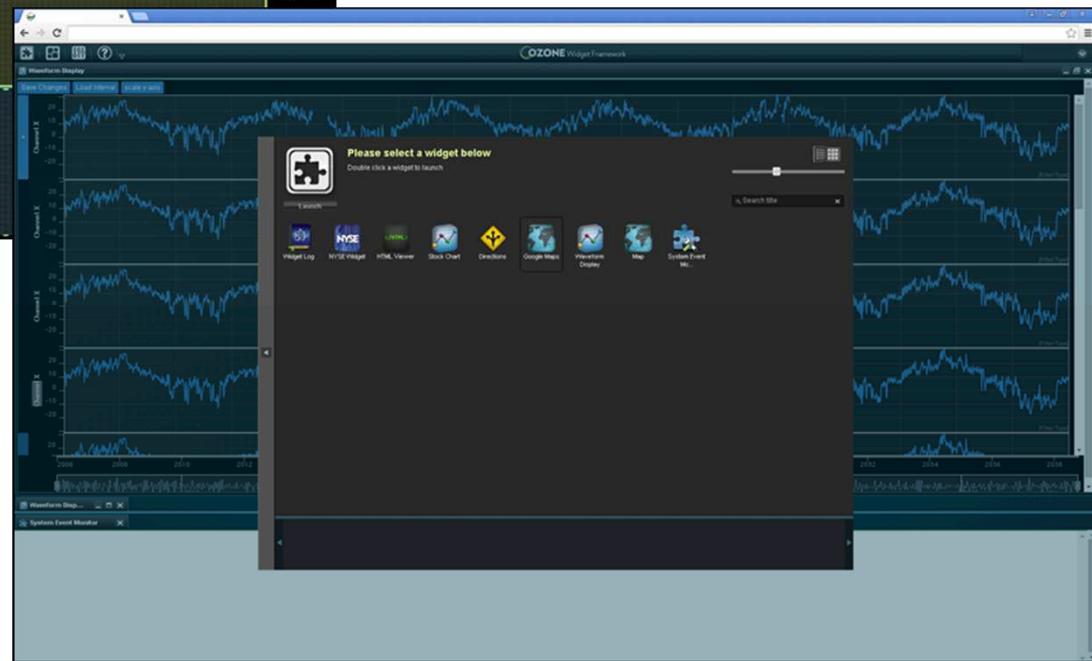


Design and save Analyst display layouts (dashboards):

- Select Widgets to display in each partition of the dashboard (drag and drop)
- E.g. waveform display, signal detection list, event detection list, map, etc.

Design and save Analyst display layouts (dashboards):

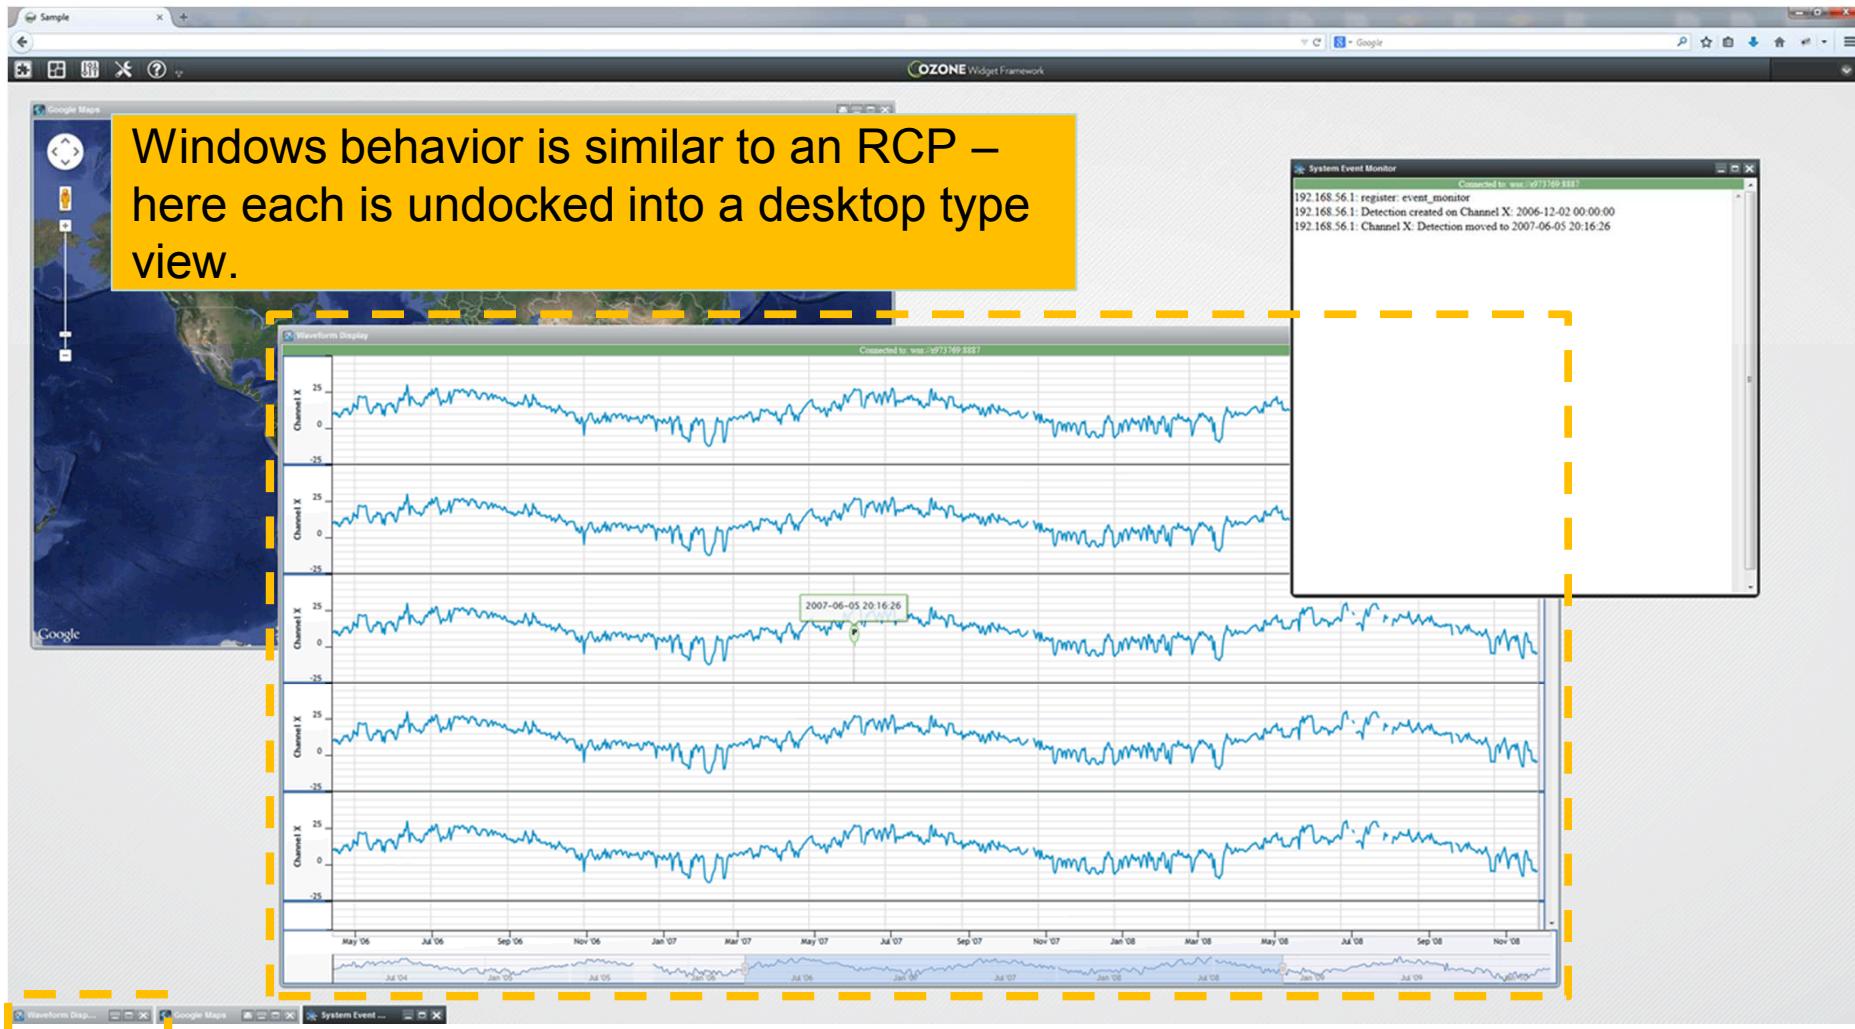
- Partition the display space
- Select *desktop, accordion, tabbed, and portal layouts for partitions*
- *Save display layouts & select upon login or configure as default*



# OWF Display Prototype – Map & Signal Detection (1 of 3)



# OWF Display Prototype – Map & Signal Detection (2 of 3)

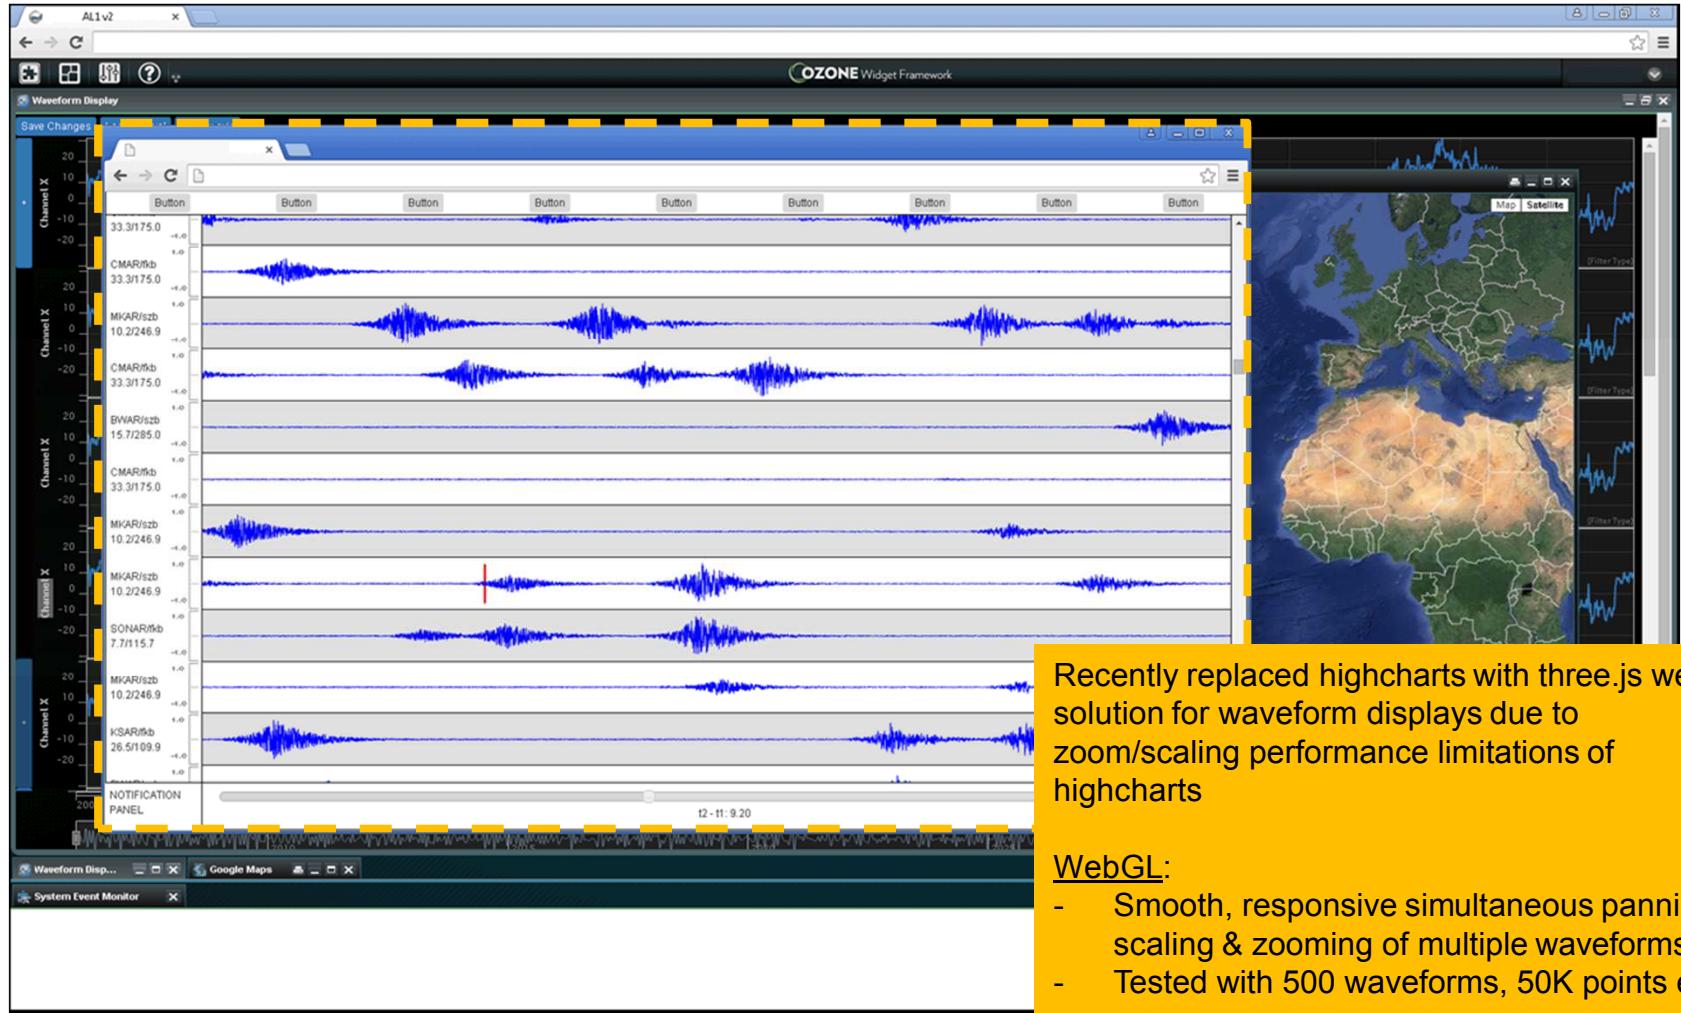


Windows behavior is similar to an RCP – here each is undocked into a desktop type view.

The screenshot shows a desktop environment with several windows. On the left, there is a 'Google Maps' window. In the center, a 'Waveform Display' window is open, showing six stacked time-series plots for 'Channel X'. The x-axis represents time from May 06 to July 09, and the y-axis ranges from -25 to 25. A green callout box highlights a specific data point on the third plot with the timestamp '2007-06-05 20:16:26'. To the right of the waveform display is a 'System Event Monitor' window, which is also undocked. It shows a list of system events connected to 'ws://192.168.56.1:8888'. The events are:

- 192.168.56.1: register: event\_monitor
- 192.168.56.1: Detection created on Channel X: 2006-12-02 00:00:00
- 192.168.56.1: Channel X. Detection moved to 2007-06-05 20:16:26

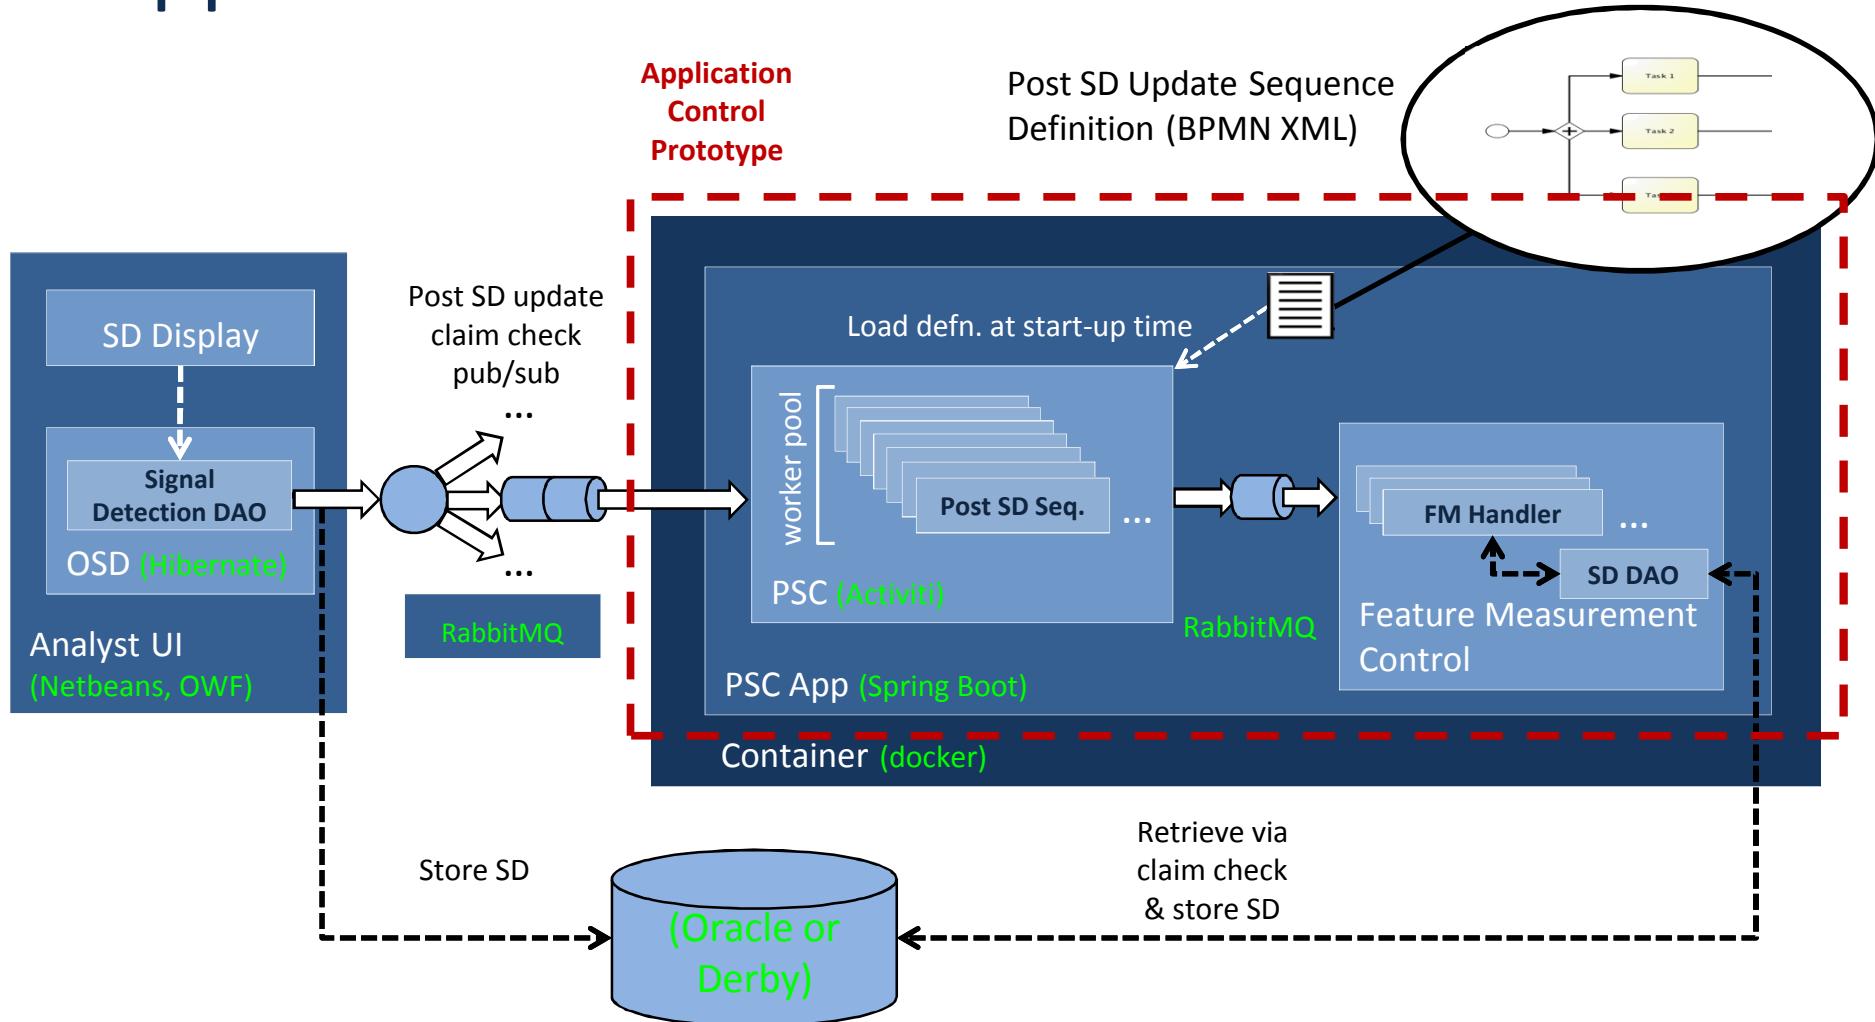
# OWF Display Prototype – WebGL Waveform Display



# Agenda

- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

# Current Prototype Status – Application Control



# Application Control

## Prototyping Status (1/4)

- Selected COTS BPMN engine (Activiti) as the basis for the PSC mechanism
  - Traded Activiti, JBoss BPMN, Spring Batch (see backup for summary comparison)
- Developed an initial PSC prototype, including:
  - Execution of mock signal detection post-processing sequences based on Analyst actions (create & update)
  - Parallel execution of mock automated signal detection processing sequences
  - Processing sequence steps delegated to domain services via RabbitMQ & REST request/response messaging interfaces

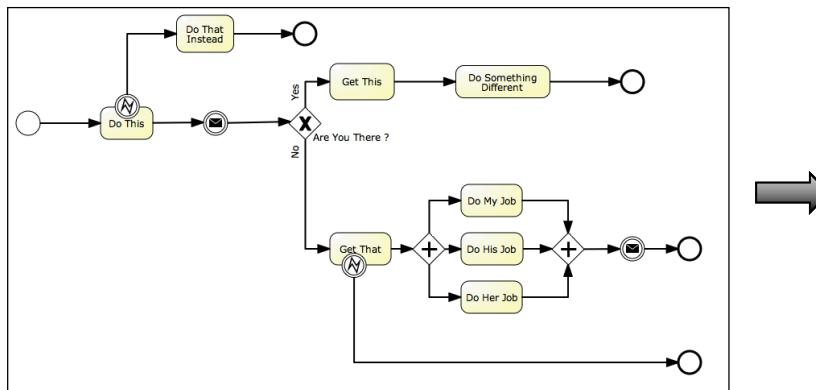
# Application Control

## Prototyping Status (2/4)

- PSC COTS evaluation focused on two primary technologies:
  - Business Process Management (BPM) Engines
    - Workflow engines providing for the definition and execution of Business Process Model and Notation (BPMN) 2.0 standard processes (BPMN standard: [www.bpmn.org](http://www.bpmn.org))
    - Prototyped Activiti BPM
    - Selected for executable architecture development
  - Java Batch Processing Engines (JSR 352)
    - “Comprehensive framework designed to enable the development of robust batch applications”
      - “Provides reusable functions that are essential in processing large volumes of records, including logging/tracing, transaction management, job processing statistics, job restart, skip, and resource management”
    - Prototyped Spring Batch

## Prototyping Status (3/4): Activiti/BPMN

- Activiti provides sequence definition and execution
  - Visual sequence definition (Activiti Designer Eclipse plugin) produces BPMN 2.0 XML processing definitions
  - Engine executes standard BPMN 2.0 XML definitions
  - Designer and engine are separable – engine works with standard BPMN 2.0 XML
- Activiti Engine is a multi-threaded runtime engine
  - Built on Spring
  - Can be embedded in any Java application (Spring or not), standalone or as a web application
- BPMN supports
  - Conditional, looped, and parallel tasks
  - Nested sequence (invoke one sequence from another sequence)
  - Timer, message driven, and rule based execution
  - REST, Camel & Mule ESB integration (e.g. initiate Activiti process from Mule, invoke Mule service from Activiti task)
  - Tasks can execute scripts and shell operations
  - Transactional sequences
  - Supports user tasks (require human intervention)



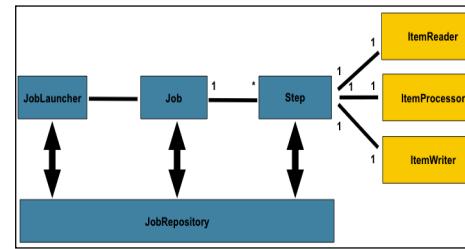
```
<?xml version="1.0" encoding="UTF-8"?>
@<definitions xmlns="http://www.omg.org/spec/BPMN/20100524/MODEL">
@  <process id="process1" name="process1">
@    <startEvent id="startevent1" name="Start"></startEvent>
@    <userTask id="usertask1" name="User Task" activiti:assignee="User">
@      <extensionElements>
@        <activiti:formProperty id="name" name="Name" type="string"></activiti:formProperty>
@      </extensionElements>
@    </userTask>
@    <endEvent id="endevent1" name="End"></endEvent>
@    <sequenceFlow id="flow1" name="to usertask" sourceRef="startevent1" targetRef="usertask1"></sequenceFlow>
@    <sequenceFlow id="flow2" name="ending" sourceRef="usertask1" targetRef="endevent1"></sequenceFlow>
@  </process>
```

## Prototyping Status (4/4): Spring Batch

- Runtime provides multi-threaded execution of batch job definitions
  - Built on Spring
  - Can be deployed standalone or as a web application
  - Transactional tasks
  - Supports remote partition-based process execution
- Job definitions support
  - Conditional, looped, and parallel tasks
  - Nested sequence (invoke one sequence from another sequence)
  - Message driven execution
  - Camel & Mule ESB integration (e.g. initiate Activiti process from Mule, invoke Mule service from Activiti task)
  - Tasks can execute scripts and shell operations
  - Supports user tasks (require human intervention)
- Limitations
  - No visual modeling support - team experienced difficulties developing complex sequences
  - No rule engine integration for rule-based job execution
  - No timer-based flow job execution
  - Limited practical documentation

```
<job id="job1" xmlns="http://Batch.Jsr352/s1">
  <split id="step1">
    <flow id="flow1" next="flow2">
      <step id="step1" next="step2" />
      <batchlet ref="MyBatchlet" />
    </step>
    <step id="step2" next="step3" />
    <batchlet ref="MyBatchlet" />
  </flow>
  <flow id="flow2">
    <step id="step3">
      <chunk reader="MyReader" processor="MyProcessor" w
        <properties>
          <property name="audit" value="true"/>
        </properties>
      </chunk>
    </step>
  </flow>
</split>
</job>
```

Job Definitions

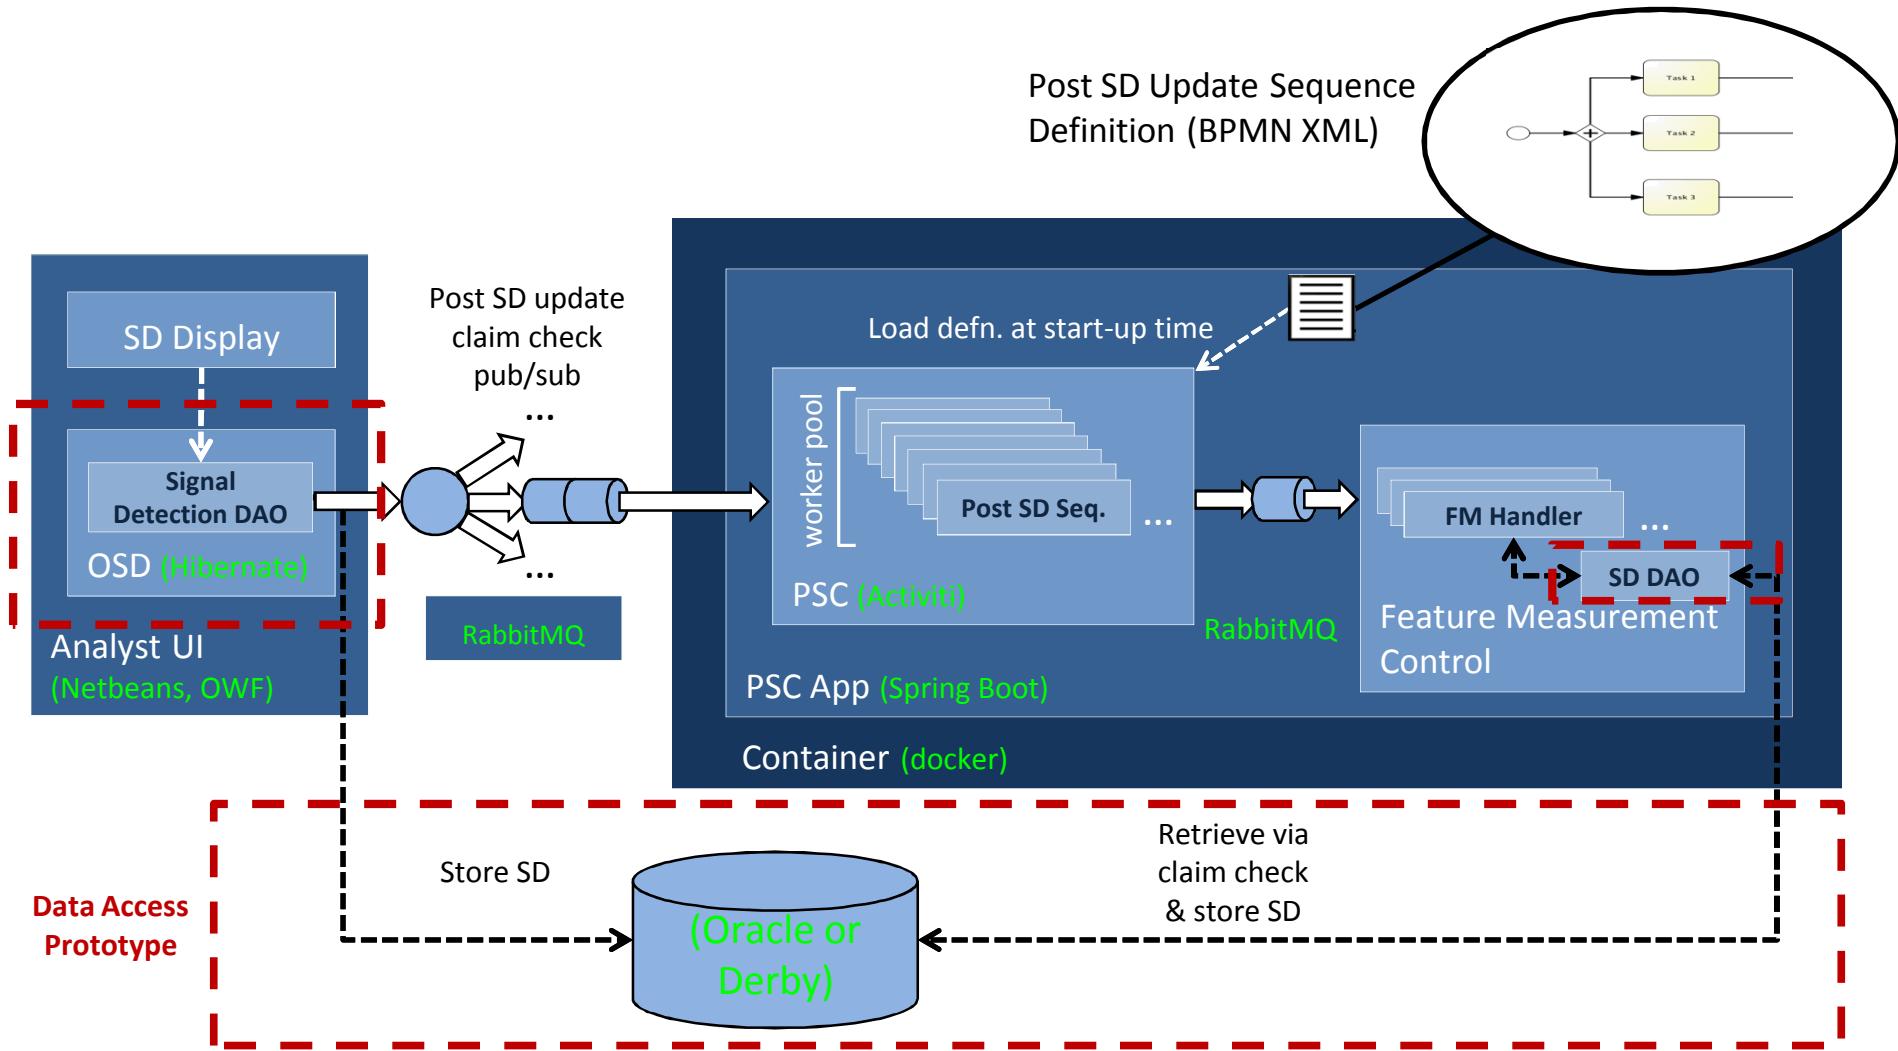


Runtime

# Agenda

- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

# Data Access Prototyping Status (1 of 2)



# Data Access Prototyping Status (2 of 2)

- Developed initial Java COI entity classes
- Developed Waveform and Signal Detection Data Accessor Objects (DAOs) providing SCRUD access to Waveform and Signal Detection entities stored in the DB
- Developed an initial OSD data distribution prototype
  - Pub/sub distribution of Signal Detection entities via RabbitMQ
    - Claim check pattern (DB reference provided via messaging, used by consumer to retrieve the entity from the DB)
    - Direct JSON serialization of entity
- Demonstrated subscription-based distribution of newly created/modified Signal Detection entities from the Analyst UI to the PSC for mock post processing

# Agenda

- Prototyping Overview
- Timeline
- Executable Architecture Prototype
  - Definition
  - Conceptual Overview
  - Assumptions
  - Current Technologies
  - Key Features
  - User Scenario
- Current Prototype Status
  - Overview
  - User Interface Prototyping Status
    - Netbeans Display Prototype
    - OWF Display Prototype
  - Application Control Prototyping Status
  - Data Access Prototyping Status
- Next Iteration Goals
- Backup

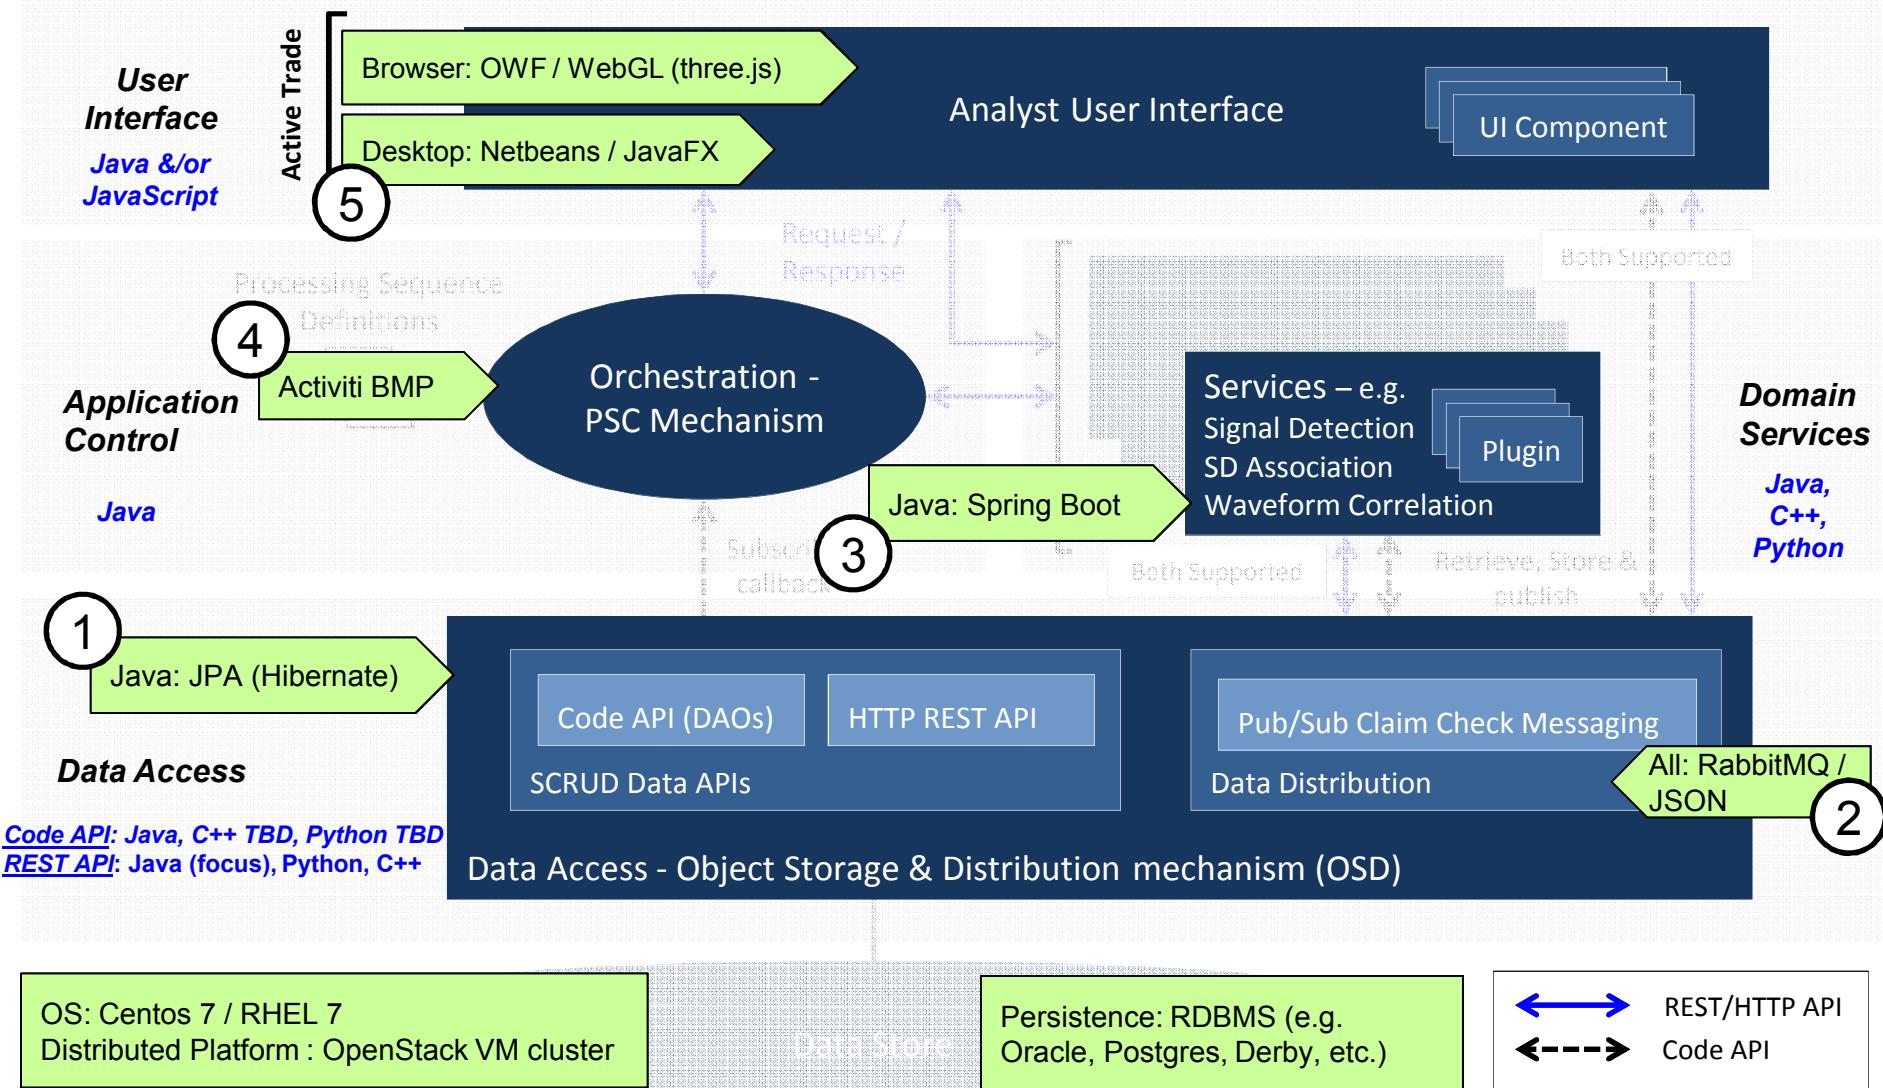
# Next Iteration Goals

- Design and develop initial prototypes for:
  - Plugin deployment and binding
  - Event history and processing parameter provenance
  - Undo/redo
- Select browser vs. RCP user interfaces
- Evaluate performance of REST-based data access
- Begin development of event analysis scenario

# BACKUP

# TECHNOLOGY EVALUATION SUMMARY

# Technology Evaluation - Overview



# 1. Technology Evaluation – Object

## Relational Mapping (data access code API)

Candidate Solution	Solution Type	Summary Assessment	
<b>Java</b>			
Hibernate	Java Object Relational Mapping (ORM) OSS	<u>Advantages</u> : Leading ORM candidate for Java. Hibernate Query Language (HQL) could provide both application and researcher level access to underlying COI objects. JPA provider.  <u>Disadvantages</u> : A dependence on HQL could introduce a tight coupling to Hibernate.	Lower database solution coupling
Open JPA	Java ORM OSS	<u>Advantages</u> : JPA provider.  <u>Disadvantages</u> : ORM features supported through embedded SQL. Not a prevalent software solution.	
Apache Cayenne	Java ORM OSS	<u>Advantages</u> : Supports Remote Object Persistence  <u>Disadvantages</u> : CayenneModeler required for mapping. Not a prevalent software solution.	
Apache Empire-DB	Java RDBMS Abstraction OSS	<u>Advantages</u> : Database interactions more easily optimized since interactions are at such a low level.  <u>Disadvantages</u> : Database abstraction layer (not an ORM). SQL-centric. Not a prevalent software solution.	
Apache Torque	Java ORM OSS	<u>Advantages</u> : Uses XML that describes the database schema, which avoids reliance on reflection.  <u>Disadvantages</u> : Requires that domain model extend Torque specific classes. Not a prevalent software solution.	Higher database solution coupling
<b>C++</b>			
ODB	C++ ORM OSS	<u>Advantages</u> : Leading ORM candidate for C++. Does not require manual entry of mapping code.  <u>Disadvantages</u> : Developed by Code Synthesis, located in South Africa. Does not provide C++ object to relational database mapping for existing DB tables.	Lower coupling
QxORM	C++ ORM OSS	<u>Advantages</u> : Supports object relational mapping with MySQL, SQLite, PostgreSQL, Oracle, and SQL Server databases.  <u>Disadvantages</u> : Market usage is unknown and documentation is limited.	Higher coupling

# 2. Technology Evaluation – Inter-process Communication (pub/sub data distribution)

Name	Standards	Language Support	Advantages	Disadvantages
RTI DDS	DDS JMS REST SOAP	C, C++ C# Java Ada	<ul style="list-style-type: none"> <li>Standards-Based</li> <li>Cross-Language Support</li> <li>Designed for low-latency, high-throughput with configurable QoS</li> <li>Flexible communication patterns &amp; configurable transports</li> <li>Open-source version available with commercial support from RTI</li> <li>Generally considered to be higher performance than brokered solutions</li> </ul>	<ul style="list-style-type: none"> <li>Open-source license is more restrictive than for other solutions</li> <li>Many features are only available in the commercial edition</li> <li>Appears to be less popular than other solutions (based on Google Trends)</li> <li>Configurable QoS introduces complexity relative to other solutions</li> <li>Past prototyping efforts have struggled with product complexity</li> </ul>
Qpid	AMQP JMS	Java C, C++ C# Ruby Perl Python	<ul style="list-style-type: none"> <li>Standards-Based</li> <li>Cross-Language Support</li> <li>Free OSS with community support</li> </ul>	<ul style="list-style-type: none"> <li>Appears to be less popular than other solutions (based on Google Trends)</li> </ul>
ActiveMQ / Apollo	AMQP STOMP REST XMPP JMS 1.1	Java C, C++ C# Ruby Perl Python	<ul style="list-style-type: none"> <li>Standards-Based</li> <li>Cross-Language Support</li> <li>Free OSS with community support</li> <li>Mature &amp; highly stable (widely used since early 2000s)</li> <li>Highly popular</li> </ul>	<ul style="list-style-type: none"> <li>Performance limitations at scale (Apollo subproject attempts to address these, but is not yet a full-featured product)</li> <li>Interest in ActiveMQ appears to be declining in recent years (based on Google trends)</li> </ul>
RabbitMQ	AMQP STOMP	Java C++ .NET Ruby Perl Python	<ul style="list-style-type: none"> <li>Standards-Based</li> <li>Cross-Language Support</li> <li>Free OSS with community support</li> <li>Commercial support available from Pivotal</li> <li>Highly popular (highest search term frequency on Google Trends)</li> <li>Favorable performance on a number of benchmarks</li> </ul>	<ul style="list-style-type: none"> <li>Broker is implemented in Erlang (not necessarily a disadvantage)</li> </ul>
ZeroMQ	None	Java C, C++ C# Ruby Perl Python	<ul style="list-style-type: none"> <li>Cross-Language Support</li> <li>Free OSS with community support</li> <li>Generally considered to be higher performance than brokered solutions</li> </ul>	<ul style="list-style-type: none"> <li>Not standards-based</li> <li>Appears to be less popular than other solutions (based on Google Trends)</li> </ul>

# 3. Technology Evaluation – Application Control

Category	Candidate Solution	Summary Assessment
<b>Enterprise Java Application Frameworks</b>	Java EE	<u>Advantages:</u> Widely-used open standards with large development community. Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. <u>Disadvantages:</u> EJB standard prohibits use of native libraries and direct thread creation, limiting design options supporting non-JVM languages.
	Spring Framework	<u>Advantages:</u> Widely-used open-source solution with large development community. Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. <u>Disadvantages:</u> Not standards-based.
<b>Stream Processors</b>	Apache Storm	<u>Advantages:</u> Open-source solution with significant industry interest. Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. Supports multiple development languages. <u>Disadvantages:</u> New offering. Not standards-based.
	Apache Samza	<u>Advantages:</u> Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. <u>Disadvantages:</u> New offering that has yet to establish significant industry interest. Not standards-based. Does not support multiple languages (Java only).
	Apache S4	<u>Advantages:</u> Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. Supports multiple development languages. <u>Disadvantages:</u> Little industry interest and development activity. Not standards-based.
<b>Enterprise Service Bus</b>	WS02 ESB	<u>Advantages:</u> Provides a robust platform for integration of heterogeneous systems via standardized messaging as part of a service-oriented architecture. <u>Disadvantages:</u> Design strengths not well aligned to the end-state modernized architecture.
<b>Complex Event Processor</b>	Esper	<u>Advantages:</u> Provides a robust platform for development of scalable, fault-tolerant, distributed processing architectures. <u>Disadvantages:</u> Specialized, query-based architecture does not fit processing needs particularly well. Not standards-based. Does not support multiple languages (Java only).

# 4. Technology Evaluation – Processing Sequence Control

Name	Standards	Advantages	Disadvantages
Activiti BPMN	BPMN 2.0	<ul style="list-style-type: none"><li>• Standards-Based</li><li>• Free OSS with community support</li><li>• Strong community, active feature development</li><li>• High-quality documentation</li><li>• Eclipse plugin integration for visual modeling</li></ul>	<ul style="list-style-type: none"><li>• Less mature than JBoss BPMN</li><li>• Commercial support by a smaller, less well known company</li></ul>
JBoss BPMN	BPMN 2.0	<ul style="list-style-type: none"><li>• Standards-Based</li><li>• Free OSS with community support</li><li>• Mature solution</li></ul>	<ul style="list-style-type: none"><li>• Decline in community development activity</li><li>• Poor documentation</li><li>• Difficult to work with</li></ul>
Spring Batch	JSR 352	<ul style="list-style-type: none"><li>• Standards-Based</li><li>• Free OSS with community support</li></ul>	<ul style="list-style-type: none"><li>• No visual modeling support - team experienced difficulties developing complex sequences</li><li>• No rule engine integration for rule-based job execution</li><li>• No timer-based flow job execution</li><li>• Limited practical documentation</li></ul>

# 5. Technology Evaluation – Desktop User Interface

Candidate Solution & Widget toolkit	Language	Summary Assessment
Netbeans / Swing	Java (RCP)	<u>Advantages</u> : Netbeans is a dominant Java UIF candidate. Swing widgets integrate alongside JavaFX code. OSGi open standard. Oracle supported. Large community. <u>Disadvantages</u> : Oracle (the company) dependence.
Eclipse / Jface (SWT)	Java (RCP)	<u>Advantages</u> : Eclipse is a dominant Java UIF candidate. OSGi open standard IBM supported. Very stable. Large community. <u>Disadvantages</u> : Eclipse learning curve is the most difficult. JFace/SWT is slightly dated compared to Swing and JavaFX2. IBM dependence.
Qt Creator / Qt	C++	<u>Advantages</u> : Qt is the leading C++ UIF candidate. GUI widgets are fast and native: strongest cross platform GUI behavior. <u>Disadvantages</u> : Not an RCP solution. Not OSGi. Smaller community than Java.
Netbeans / JavaFX2	Java (RCP)	<u>Advantages</u> : Netbeans is the leading Java UIF candidate. JavaFX2 has most modern Java GUI elements. OSGi open standard. Oracle supported. Large community. <u>Disadvantages</u> : JavaFX2 2D plotting package is beautiful but has serious scaling issues. Oracle dependence.
NA / wxWidgets	C++	<u>Advantages</u> : Native mode widget toolkit, also contains inter-process communication layer <u>Disadvantages</u> : Not an RCP solution or a UIF - mainly a standalone widget toolkit. Smaller community.
NA / XUL	XML & Java	<u>Advantages</u> : XML markup language for GUI construction. Quick study for web designers. <u>Disadvantages</u> : Not an RCP solution or a UIF - mainly a standalone widget toolkit. Not a prevalent solution.

# Technology Evaluation – In-Memory Caching/Data Grid

Selection pending determination of need for caching

Name	Client Language Support	Advantages	Disadvantages
JCS	Java	<ul style="list-style-type: none"> <li>• Cross-Language Support</li> <li>• Free OSS with community support</li> </ul>	<ul style="list-style-type: none"> <li>• Java only (no cross-language support)</li> <li>• Appears to be less widely used/popular than other solutions (e.g. Redis, memcached)</li> <li>• It is not clear whether commercial support is available</li> <li>• Limited feature set relative to other solutions surveyed</li> <li>• Does not support partitioning (only replication)</li> </ul>
memcached	C, C++ Java, Python Ruby Perl C#	<ul style="list-style-type: none"> <li>• Well established and mature</li> <li>• Widely used highly popular</li> <li>• Cross-Language Support</li> <li>• Free OSS with community support</li> <li>• Commercial support available</li> </ul>	<ul style="list-style-type: none"> <li>• Popularity appears to be declining (based on Google Trends)</li> </ul>
EHCache	Java C++ &C# (commercial version)	<ul style="list-style-type: none"> <li>• Cross-Language Support</li> <li>• Free OSS version available</li> <li>• Commercial support available from Terracotta</li> <li>• Strong feature set, including partitioning, replication, transactions, security, etc.</li> </ul>	<ul style="list-style-type: none"> <li>• Many features are only available in the commercial edition</li> <li>• Limited cross-language support (and only in the commercial edition)</li> <li>• Appears to be less widely used/popular than other solutions (e.g. Redis, memcached)</li> <li>• Popularity appears to be declining (based on Google Trends)</li> </ul>
Infinispan	C++ Java Python Ruby C#	<ul style="list-style-type: none"> <li>• Cross-Language Support</li> <li>• Free OSS with community support</li> <li>• Commercial support available from JBoss</li> <li>• Strong feature set, including partitioning, replication, transactions, security, etc.</li> </ul>	<ul style="list-style-type: none"> <li>• Appears to be less widely used/popular than other solutions (e.g. Redis, memcached)</li> </ul>
Redis	C, C++ Java Perl Python Ruby C# Closure Scala	<ul style="list-style-type: none"> <li>• Widely used highly popular</li> <li>• Broad cross-Language Support</li> <li>• Free OSS with community support</li> <li>• Commercial support available from Pivotal</li> <li>• Strong feature set, including partitioning, replication, transactions, etc.</li> </ul>	<ul style="list-style-type: none"> <li>• Limited built-in security features</li> </ul>
Hazelcast	Java C++ &C# (commercial version)	<ul style="list-style-type: none"> <li>• Cross-Language Support</li> <li>• Commercial support available from Hazelcast</li> <li>• Strong feature set, including partitioning, replication, transactions, security, etc.</li> </ul>	<ul style="list-style-type: none"> <li>• Many features are only available in the commercial edition</li> <li>• Limited cross-language support (and only in the commercial edition)</li> <li>• Appears to be less widely used/popular than other solutions (e.g. Redis, memcached)</li> </ul>

# DOCKER CONTAINER TECHNOLOGY OVERVIEW

# Agenda

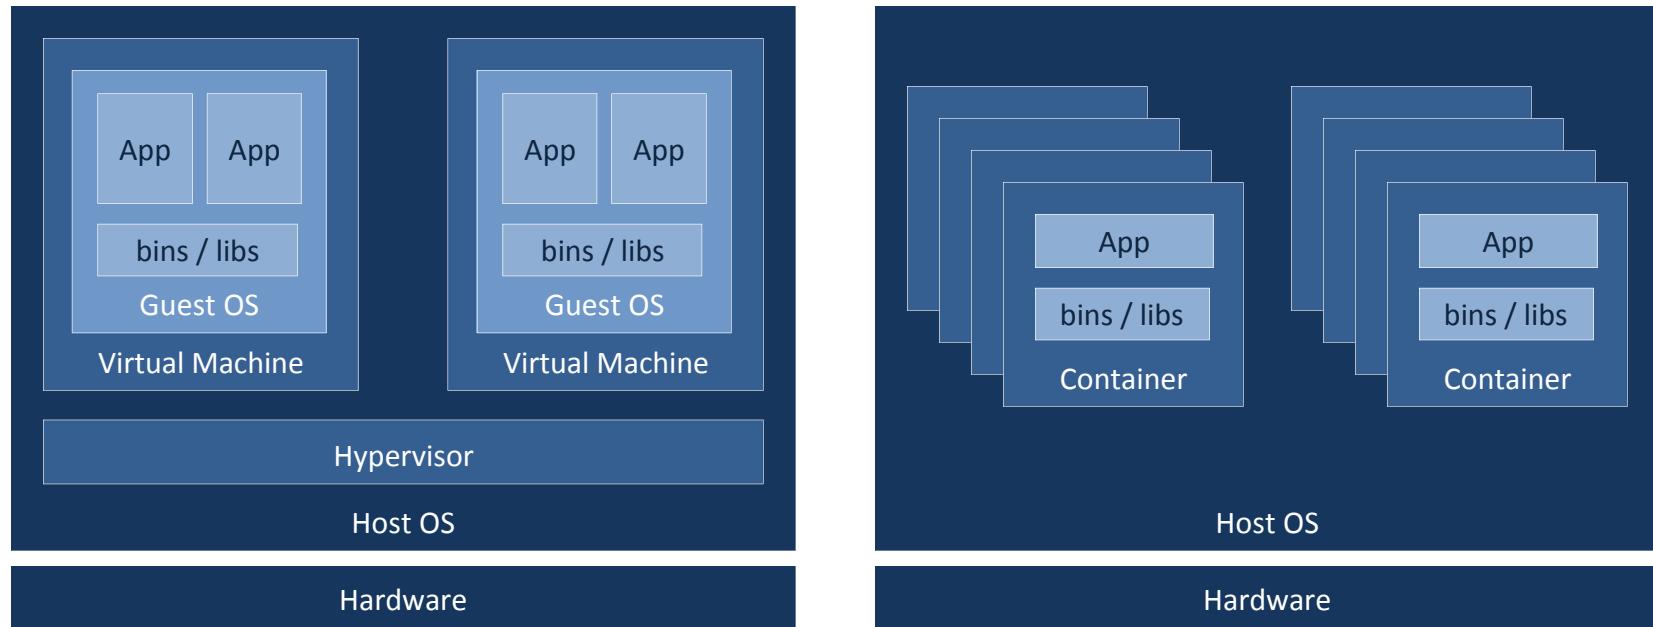
- What are Linux Containers?
- Containers vs. Virtual Machines
- Why Containers?
- Container Ecosystem
- What is Docker?
- Container Use Cases
- Development Lifecycle Concepts
- Integration Concepts
- Orchestration Concepts

# What Are Linux Containers?

- Lightweight OS-level virtualization technology
- Multiple, isolated systems (containers) run on a single Linux host, sharing the underlying kernel
  - Control Groups provide resource isolation (CPU, memory, block I/O, network, etc.)
  - Namespaces isolate applications' view of the OS environment (processes, networking, file system, etc.)
- Alternative paradigm to virtual machines
  - Containers offer greatly reduced start-up time & resource utilization, as well as near-native speed
    - CPU performance – native
    - Memory – Very small overhead
    - Network – Very small overhead (can be optimized to near zero)
  - Individual applications & services are typically deployed in separate containers (100s-1000s of containers per node is common)

# Containers vs. Virtual Machines

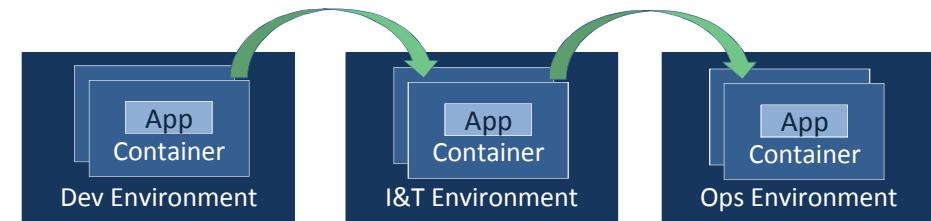
- Unlike virtual machines, containers provide separate virtual OS environments that share the underlying host OS directly, without the need for separate guest OS instances or hypervisor software



# Why Containers?

## Common Use Cases

- Isolate individual *applications*
- Constrain application resource utilization
  - CPU, memory, network & storage
- Manage heterogeneous application dependencies across complex systems
  - Multiple software stacks, third party tools & versions
  - Multiple Linux flavors & versions
- Provide orchestration of applications across clustered system deployments
- Support application scalability and fault tolerance
- Provide a consistent application runtime for development through operational deployment



# Container Ecosystem

- A thriving community of open-source container technologies has emerged in the last few years supporting the development and deployment of container-based systems
- Large impact on industry with significant backing and community interest
  - Widely used in modern PaaS solutions – e.g. Amazon EC2, DigitalOcean, Google Compute Engine, Microsoft Azure, OpenStack, QEMU/KVM, Vagrant and Vmware
  - Widely used in corporate IT infrastructure - **Google's infrastructure runs nearly entirely on containers (2 billion+ provisioned per week)**
- **Docker has emerged as the dominant container project around which the majority of these technologies is currently developed**

Orchestration	Swarm Compose Machine	Fleet	Kubernetes	Mesos	
Container Management	Docker	Rocket	LXC		
Host OS	RedHat Linux	CentOS	CoreOS	Ubuntu	Project Atomic

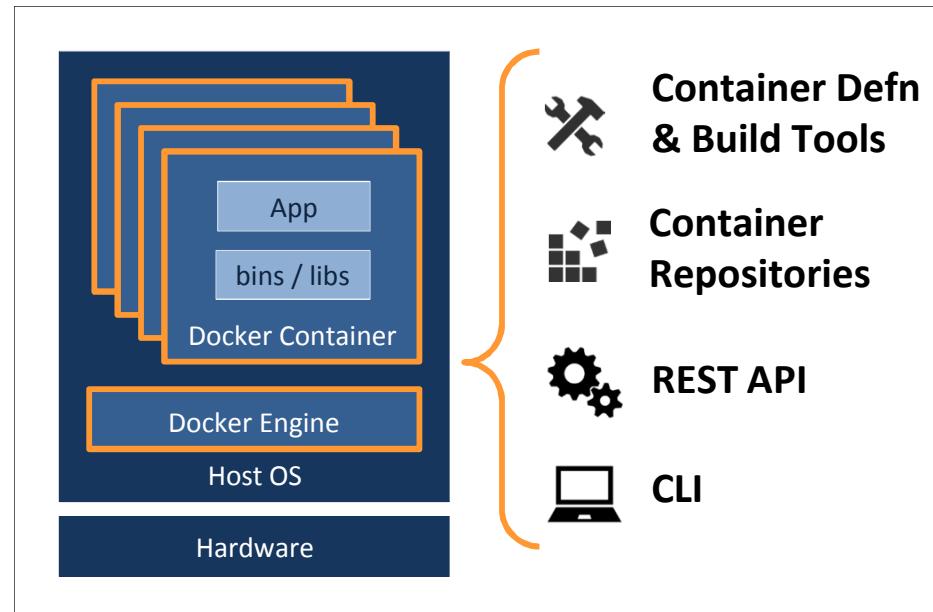
Load-balancing, scalability & high availability for distributed, container-based systems

Definition, build, & deployment of containerized apps

OS shared across containers

# What is Docker?

- Open-source project providing container abstraction and management tooling
- Container Definition & Build
  - Build containers from pre-defined images
  - Define images declaratively by specifying new layers on top of existing images, from the base OS up
- Container Repositories
  - Manage container images for projects & communities using shared repositories
  - Publish and consume public images with docker hub (the definitive public repository)
- REST API & CLI
  - Define, build, deploy & manage containers via CLI and RESTful service interfaces



# Potential US NDC/IDC Container Use Cases

- Development Lifecycle
  - Provide a common software platform for application developers across environments, from development, to I&T and operations
- Integration
  - Enable integration of heterogeneous applications developed by multiple organizations on a common platform
    - Isolation, resource constraints, support for multiple software languages & stacks, third-party COTS, versions, OS versions, etc. as needed
- Orchestration
  - Provide software stack-agnostic, scalable, fault-tolerant orchestration of applications and services in a distributed environment
- Deployment
  - Support for deployment to bare metal, VM and cloud platforms

# Development Lifecycle Concepts

- A common application platform is provided to US NDC/IDC contributors as a set of base container images
  - Includes OS, third-party tools & libraries, core framework software
  - Provided through a shared container registry
- Applications are developed using the base platform containers and are packaged as new container images layered onto the base platform
  - Application containers are delivered back to the shared container registry
  - Continuous build & integration includes automated build and test of containerized applications & services
- Application containers are deployed into the US NDC/IDC environments as part of the development lifecycle, including development, I&T and operations

# Integration Concepts

- The system integrator organization maintain the configuration needed to deploy and manage the set of application containers comprising the system
  - System cluster definition
  - Application definitions
  - Application resource constraints
  - Network and persistent storage interface management
  - High-availability & scalability policies
- Multiple libraries, third-party tools, versions, etc. are encapsulated within individual application containers where they are required