Modeling and Predicting HLA Federation Performance

Stephen Kolek, Steve Boswell, Harry Wolfson

MIT Lincoln Laboratory

Distributed Systems Performance Modeling Office (DSPMO)

00F-SIW-085

18 September 2000

Overview

- The Federation Performance Modeling Consortium
- What are We Trying to Accomplish?
- Modeling Approach
- Tools Required
- Progress to Date

What Federations Do

- **HLA Federations can have different performance goals**
  - **Time-managed, fast-as-possible**
  - **Wall clock time (time "happens")**
  - **Time-managed, human-in-the-loop**
  - **Minimum latency, hardware-in-the-loop**
- **Federation divvies up the virtual world, considering**
  - Functionality**: types of simulated objects, sources, granularity, etc.**
  - Modularity**: decomposition, reusability, potential extensions**
  - host performance**: computation throughput, available memory**
  - network performance**: bandwidth and latency between federates**
- **Federates exchange object attributes, interactions via RTI**
  - **(Some) RTIs attempt to minimize traffic over the network with data distribution management, message aggregation, etc.**
  - **RTIs may offer other services such as time management, reliable transport**

Predicting Federation Performance

- **Goals**
  - Reduce cost of Federation development and debugging
  - Prediction by extrapolation on observables using models
- **Deliverables**
  - Predictive modeling tools and user guides
  - Initial predictive models for a specific RTI (1.3NG)
  - Catalog of observable features and metrics
    - Developed during model prototyping; useful for predicting performance of similar Federations
  - Tools to extend the catalog

An Iterative Process

- Goals
- **Research**
  - **Conceptualize models, define observables**
    - For some specific Federates and specific RTIs
  - **Statistical nature of models and observables**
    - Can't hope to know all details
    - Model salient behaviors and their variability
- **Modeling**
  - **Build prototype predictive models**
- **Measurement and Analysis**
  - **Are measured observables sufficient, efficient, etc?**
  - **How accurate, reliable are predictive models?**
- Deliverables

How Is This Effort Different?

- **Benefits HLA Community at Large**
  - **Relies on input from Consortium members**
  - **Models and tools must be relevant, useful, and cost effective to broad constituency**
- **Focused on Boundary Behavior**
  - **Exchanges between the Federates and the RTI**
  - **Not dependent on detailed internal models of Federates**
  - **Modeling effects of Federate activity**
- **Quantitative Models for Performance Prediction**
  - **Getting beyond performance of two or three test Federates with "evenly" balanced work loads**
  - **Tools to aid Federation development**
    - Is my Federation working as predicted at this stage?

Overview

- **The Federation Performance Modeling Consortium**
- **What are We Trying to Accomplish?**
- **Modeling Approach**
  - Performance Feasibility Limits
  - Nominal Performance
- **Tools Required**
- **Progress to Date**

Task 1: Feasibility Limits

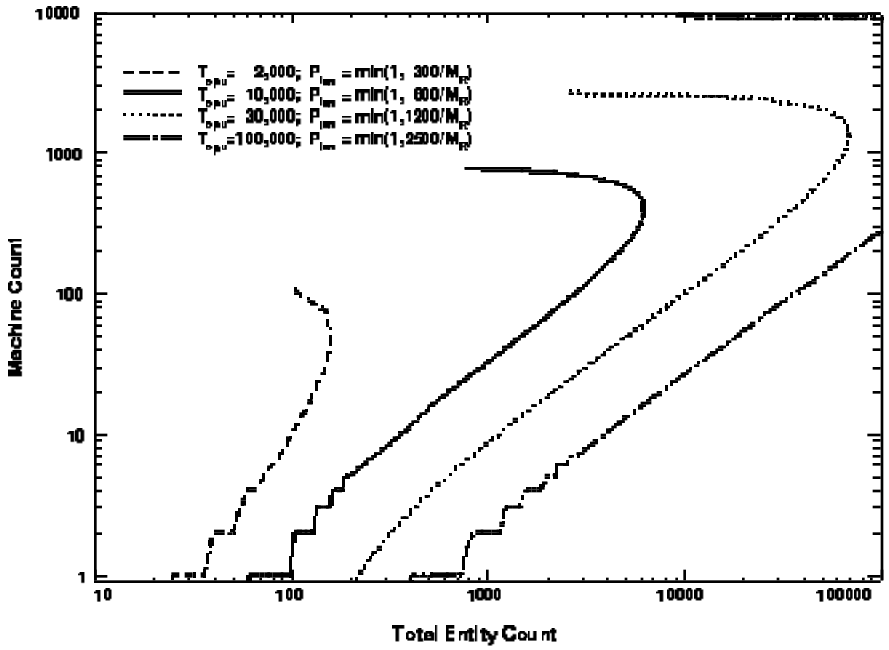

- More federates, more objects, demand more processing resources
- Limited resources per host requires greater distribution of processing
- Greater distribution of processing requires more communication processing
- At some point, all resources end up being consumed by communication processing!

Predict whether a proposed Federation will overwhelm available resources.

Host Resource Constraints

- Each host has resource limits
  - CPU cycles/sec
  - Random Access Memory (RAM)
- Per-host resource costs include
  - OS overhead to maintain Federate, RTI, other apps
  - Cost of simulating objects
  - Cost of processing updates and interactions via the RTI
- Available RAM affects CPU load
  - Less RAM necessitates more swapping to/from disk
- Violation of resource limits is catastrophic
  - Dropped messages, unacceptable latency, time stalls, etc.

Feasibility Limit Model

Each host must satisfy

$$\left(\begin{array}{c}\text{Simulation}\\ \text{Overhead}\end{array}\right)+\left(\begin{array}{c}\text{Simulation}\\ \text{Processing}\end{array}\right)+\left(\begin{array}{c}\text{Message}\\ \text{Processing}\end{array}\right)\leq T_{CPU}$$

- Intrinsic costs of application, RTI, OS support
- RTI management services costs per Federation layout
- Maintenance costs of objects and attributes

- Task assignment costs for local objects
- Task execution costs for local objects
- Response to Interactions
- Generation of Interactions

- Per Message Costs of RTI Services

Limit Model is a set of constraint relationships that must be satisfied for the most stressing activity distribution.

Task 2: Nominal Behavior Modeling

- **Predict how well the Federation will satisfy performance requirements**
  - Minimum latency, hardware-in-the-loop
  - Time-managed, human-in-the-loop
  - Time-managed, fast-as-possible
  - Wall clock time (time "happens", 3 hours or 3 days)
- **Requires models of message statistics**
  - Message source statistics
  - Effects of RTI and network on messages

Predictive Model Components

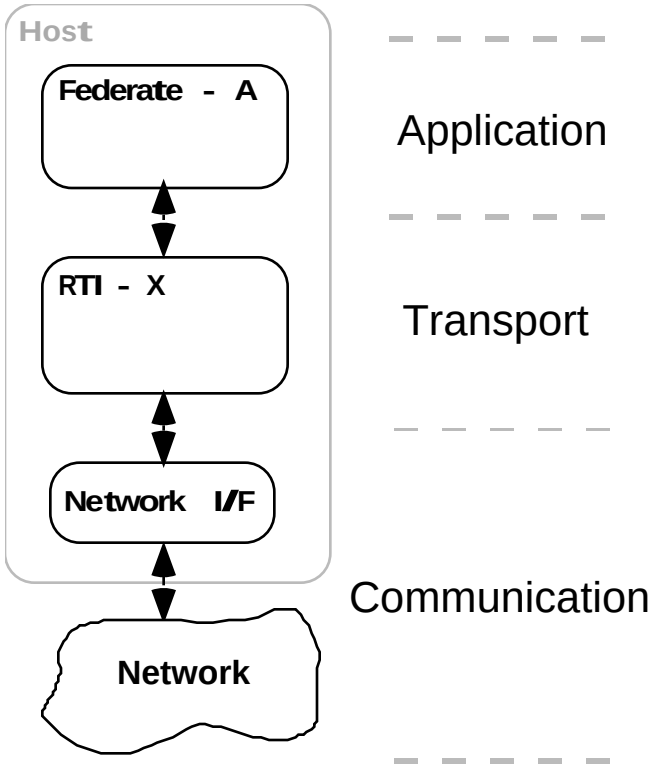

- **Each Federate emits and collects messages**
  - **messages are statistically distributed in size, time of emission, destination**
- **Transport model (RTI-specific)**
  - **Forwards messages based on subscriptions**
  - **Alters message distributions**
    - Time managers, bundlers, and reliable transport alter time distributions
    - DDM alter destination distributions
- **Network model**

RTI Behavior Model

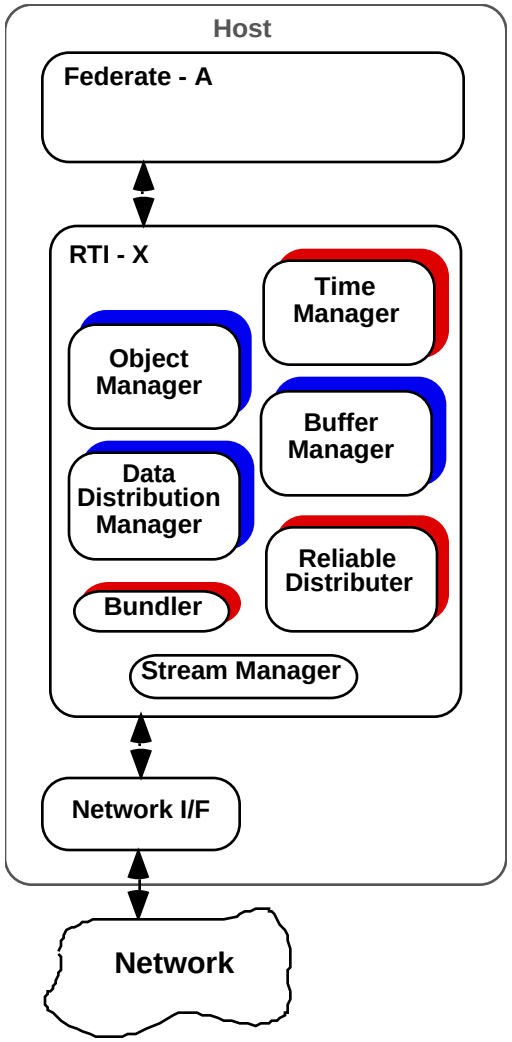

- **RTI's "message relevance filtering" saves Federate CPU cycles, but costs message delay and memory consumption**
  - Time Manager delays messages
  - Buffer Manager consumes memory to store pending messages
- **Some of these effects have been modeled by others**
- **This study will tie together existing models and create new models as needed**

Overview

- The Federation Performance Modeling Consortium
- What are We Trying to Accomplish?
- Modeling Approach
- Tools Required
- Progress to Date

RTI Wrapper

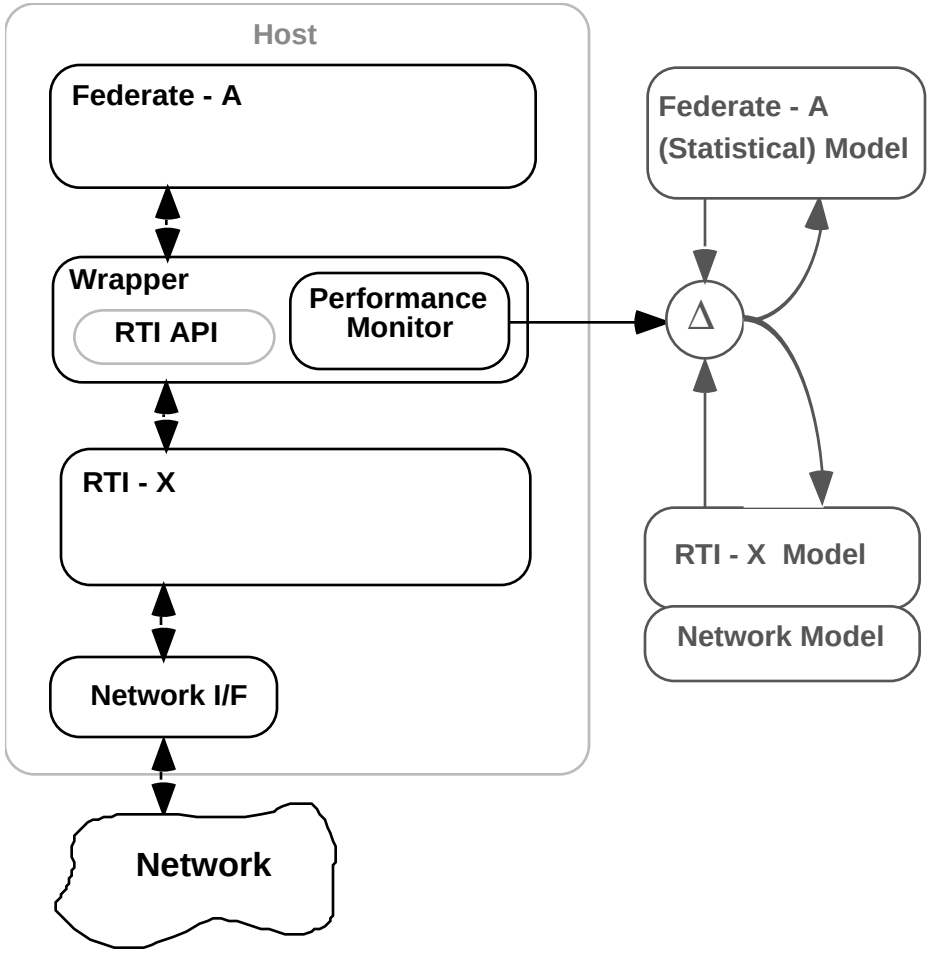

- **Generic interface layer between any Federate and any RTI**
- **Produces statistical snapshots**
  - Federate and RTI exchanges
  - Federate and RTI resource consumption (memory, CPU)
- **Supports real-time performance monitor**
- **Provides "ground truth" data for development and verification of performance models**

RTI Wrapper

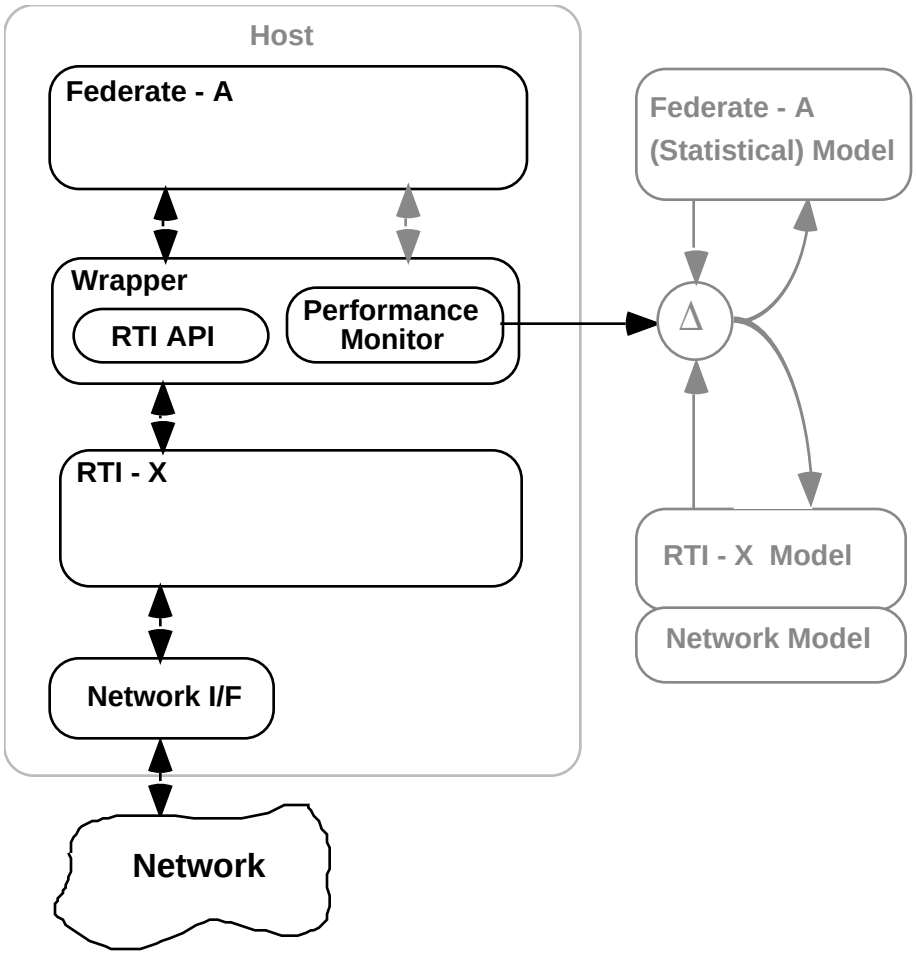

- Federate calls the Wrapper just like the RTI (API is identical)
- Instrumentation capabilities are open-ended:
  - extensible catalog of measurables
  - a Federate can plug in customized measurement functions
- Run-time configurable (what is measured, when, and how)

Run-Time Configuration Control

Which API methods

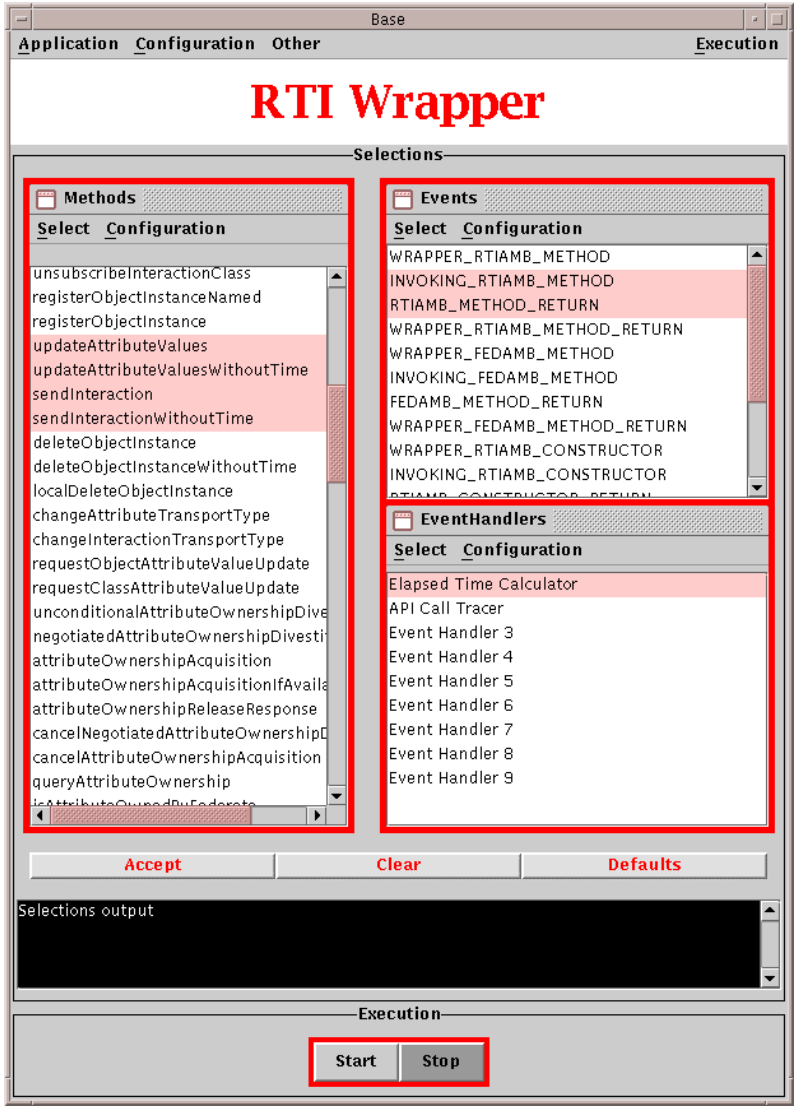

Which points in execution flow

Which performance measurements

Configuration Panel: Shortcuts

Quick-keys

Save run-time configurations

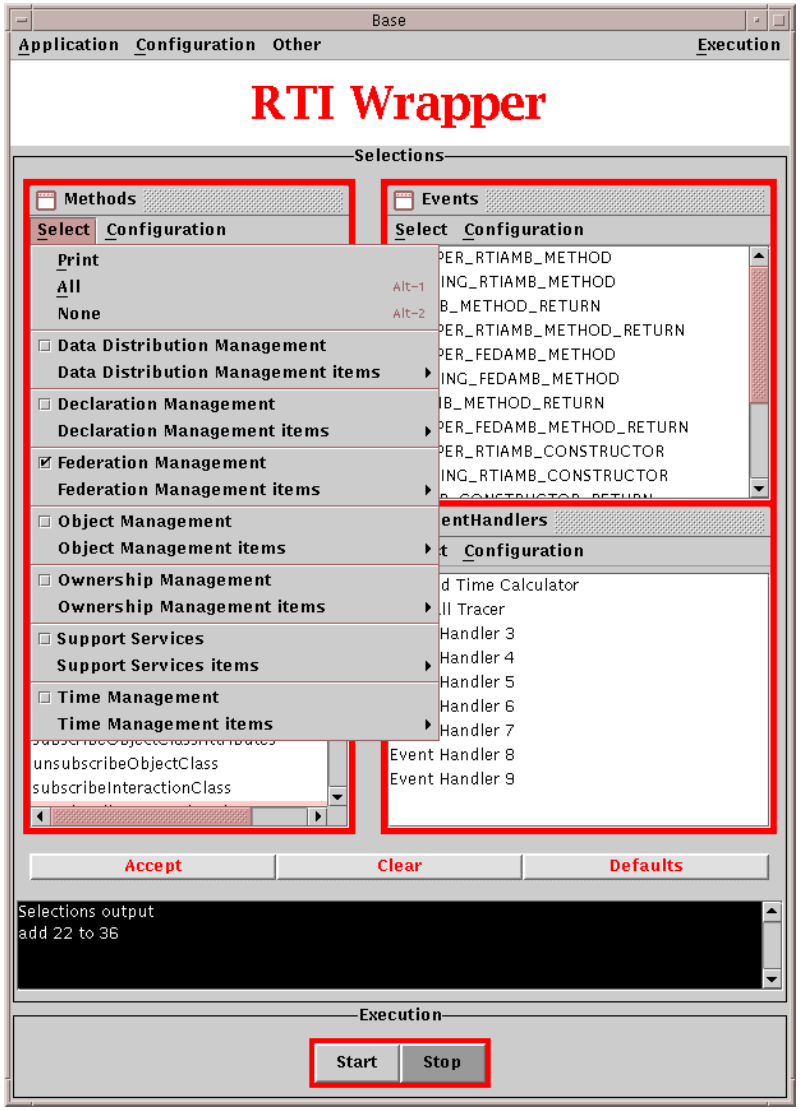

Note:

- **GUI requires no Federate code**
- **API controls also available to Federate**

Wrapper: Progress to Date

- Built a framework in which instrumentation can be developed
- Supplied a few native instrumentation functions
- Framework allows Federation developers to customize and control instrumentation capabilities
- Framework is usable with virtually no Federate code changes
- Adapted example applications to demonstrate use of the Wrapper (HelloWorld, Benchmarks)
- Initial release anticipated 15 October

Wrapper: Future (with Consortium)

- Extend maturity and breadth of instrumentation functions
- Port to additional platforms
- Additional testing
- Develop protocols for formatting and reporting output
- Develop mechanisms for collaboration, control, and information sharing across a Federation execution

Network Sounder

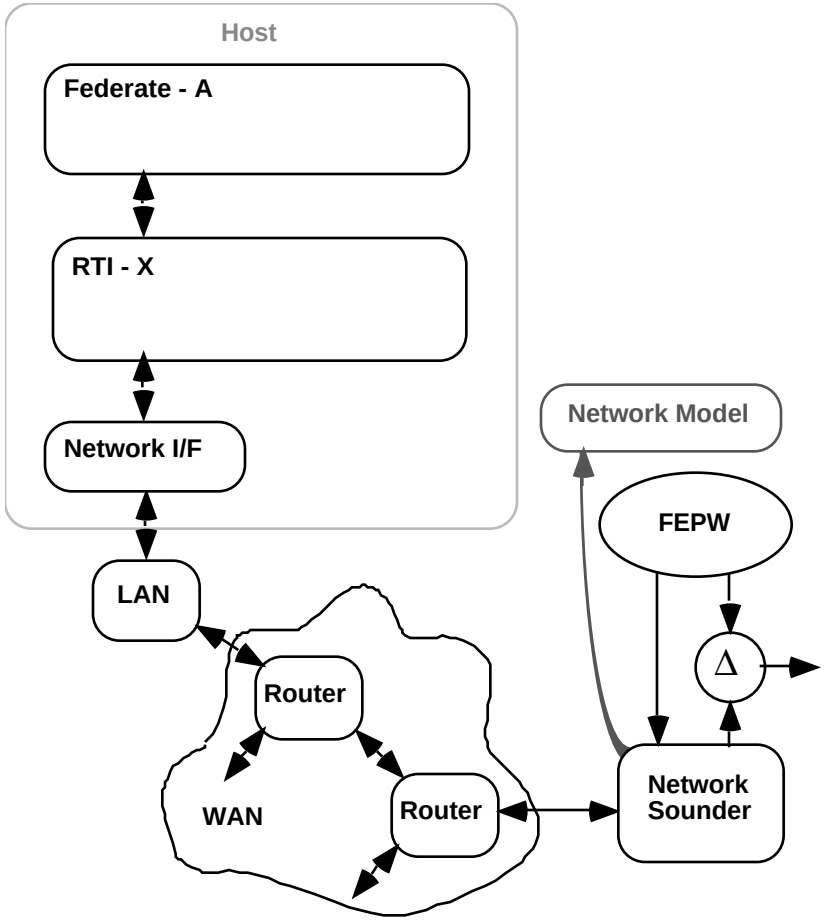

- **Builds Federation connectivity model**
  - Initialized by Federation host table (from FEPW, etc.)
- **Capabilities**
  - Verify required host connectivity
  - Measure and quantify network performance
  - Discover (unanticipated?) non-Federation related traffic
  - Aid debugging
- **Assertion**
  - IP networks will require the most complex sounder

NetSounder: Near Term Capability

- **Verify basic network connectivity**
  - **Import Federate host IP info from FEPW DIF file, or alternate source**
  - **Stand Alone Application** (RTI generic)
    - DNS lookup
    - ICMP / TCP Ping
    - Traceroute
    - Telnet (useful to aid in debugging)
  - **Distributed / Coordinated Application** (OS / RTI specific)
    - TCP / UDP connections
    - Multicast verification
    - Reliable and Best Effort transport
- **Initial Release anticipated 15 Oct 2000**

NetSounder: Future Capability

- **Measure network performance**
  - Identify network topology (routers, etc)
  - Intrinsic latency between Federate host platforms
  - Sustainable throughput (bandwidth)
- **Identify network loads**
  - Verify that Federate generates anticipated levels of network traffic
  - Identify excessive network traffic not associated with Federation
- **Network performance statistics used to develop Federation predictive models**

NetSounder: Host Table

- **Federate host table read from FEPW or alternate source**

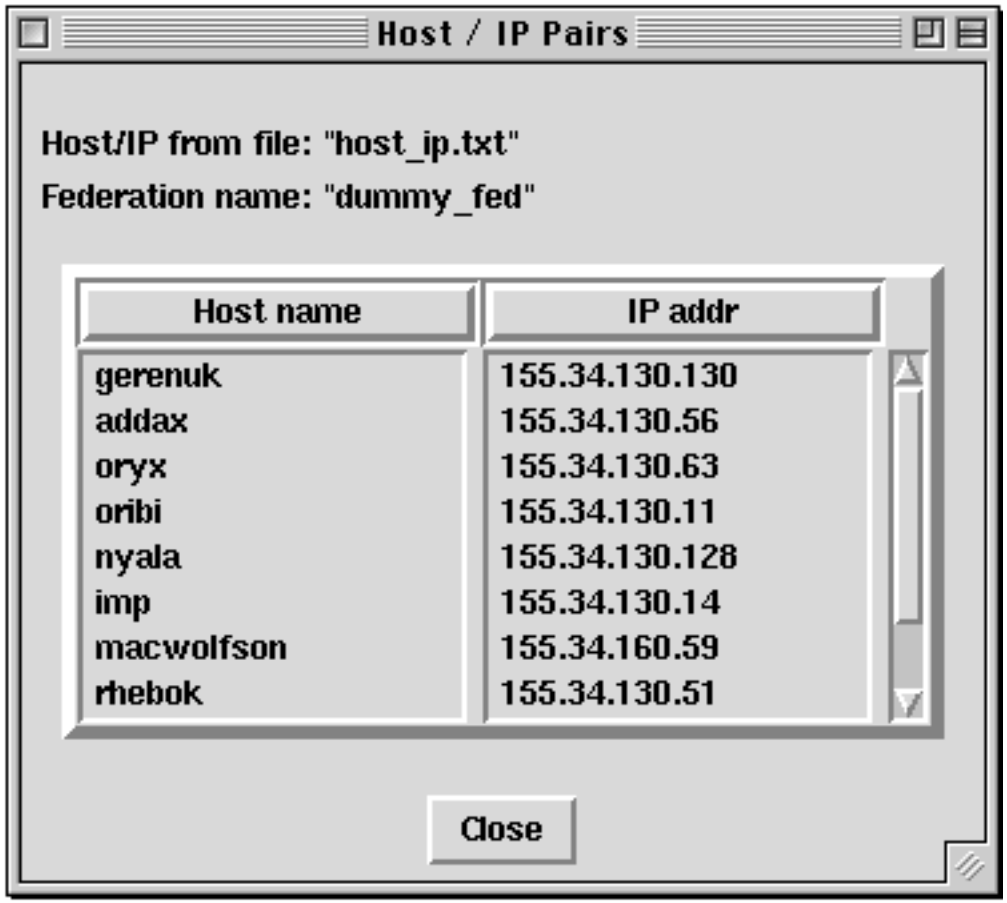

NetSounder: Network Tests

- **Verify basic network connectivity**

Network Connectivity Tests

Host/IP from file: "host_ip.txt"
Federation name: "dummy_fed"
Network Connectivity tests run from host: "gerenuk"

Host name	IP addr	DNS Lookup	OS ICMP ping	Perl TCP ping	Traceroute	Telnet Probe
gerenuk	155.34.130.130	pass	n/a	n/a	n/a	pass
addax	155.34.130.56	pass	n/a	n/a	n/a	pass
oryx	155.34.130.63	pass	n/a	n/a	n/a	pass
oribi	155.34.130.11	pass	n/a	n/a	n/a	pass
nyala	155.34.130.128	pass	n/a	n/a	n/a	pass
imp	155.34.130.14	FAIL	n/a	n/a	n/a	FAIL
macwolfson	155.34.160.59	pass	n/a	n/a	n/a	FAIL
rhebok	155.34.130.51	pass	n/a	n/a	n/a	pass

Close

NetSounder: Diagnose Errors

- **Diagnose host connectivity problems**

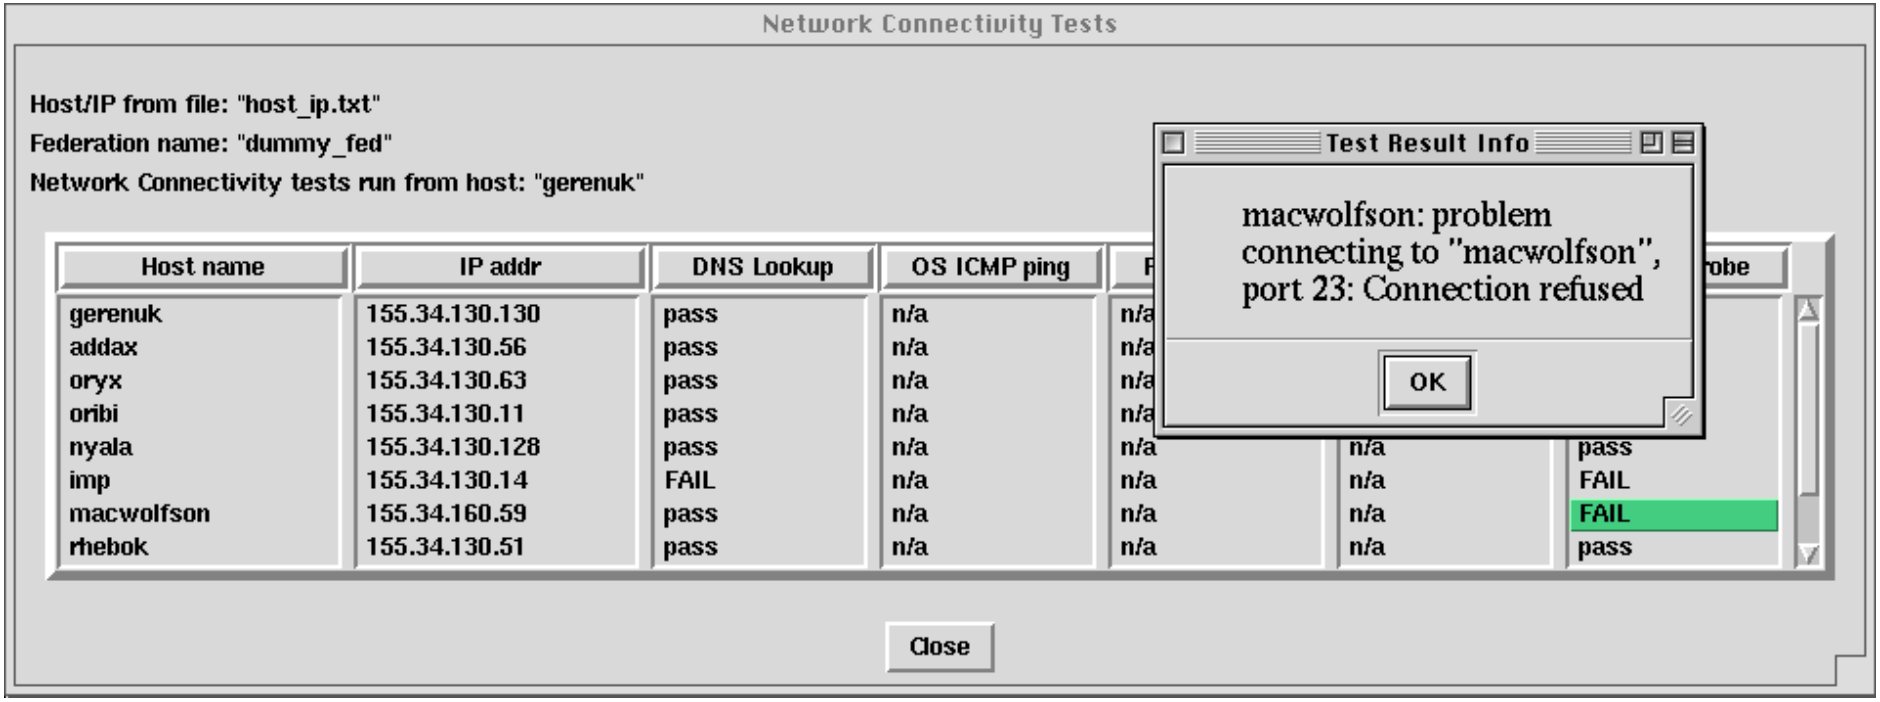

Overview

- The Federation Performance Modeling Consortium
- What are We Trying to Accomplish?
- Modeling Approach
- Tools Required
- Progress to Date

Progress to Date

- √ Solicit consortium feedback on concepts
- √ Search for existing tools and models
- √ Develop initial suite of tools
- √ Solicit SIW community feedback on concepts
- √ Establish mechanism for tool distribution

http://DSPMO.LL.MIT.EDU

info@DSPMO.LL.MIT.EDU

DSPMO Web Site

http://DSPMO.LL.MIT.EDU **info@DSPMO.LL.MIT.EDU**

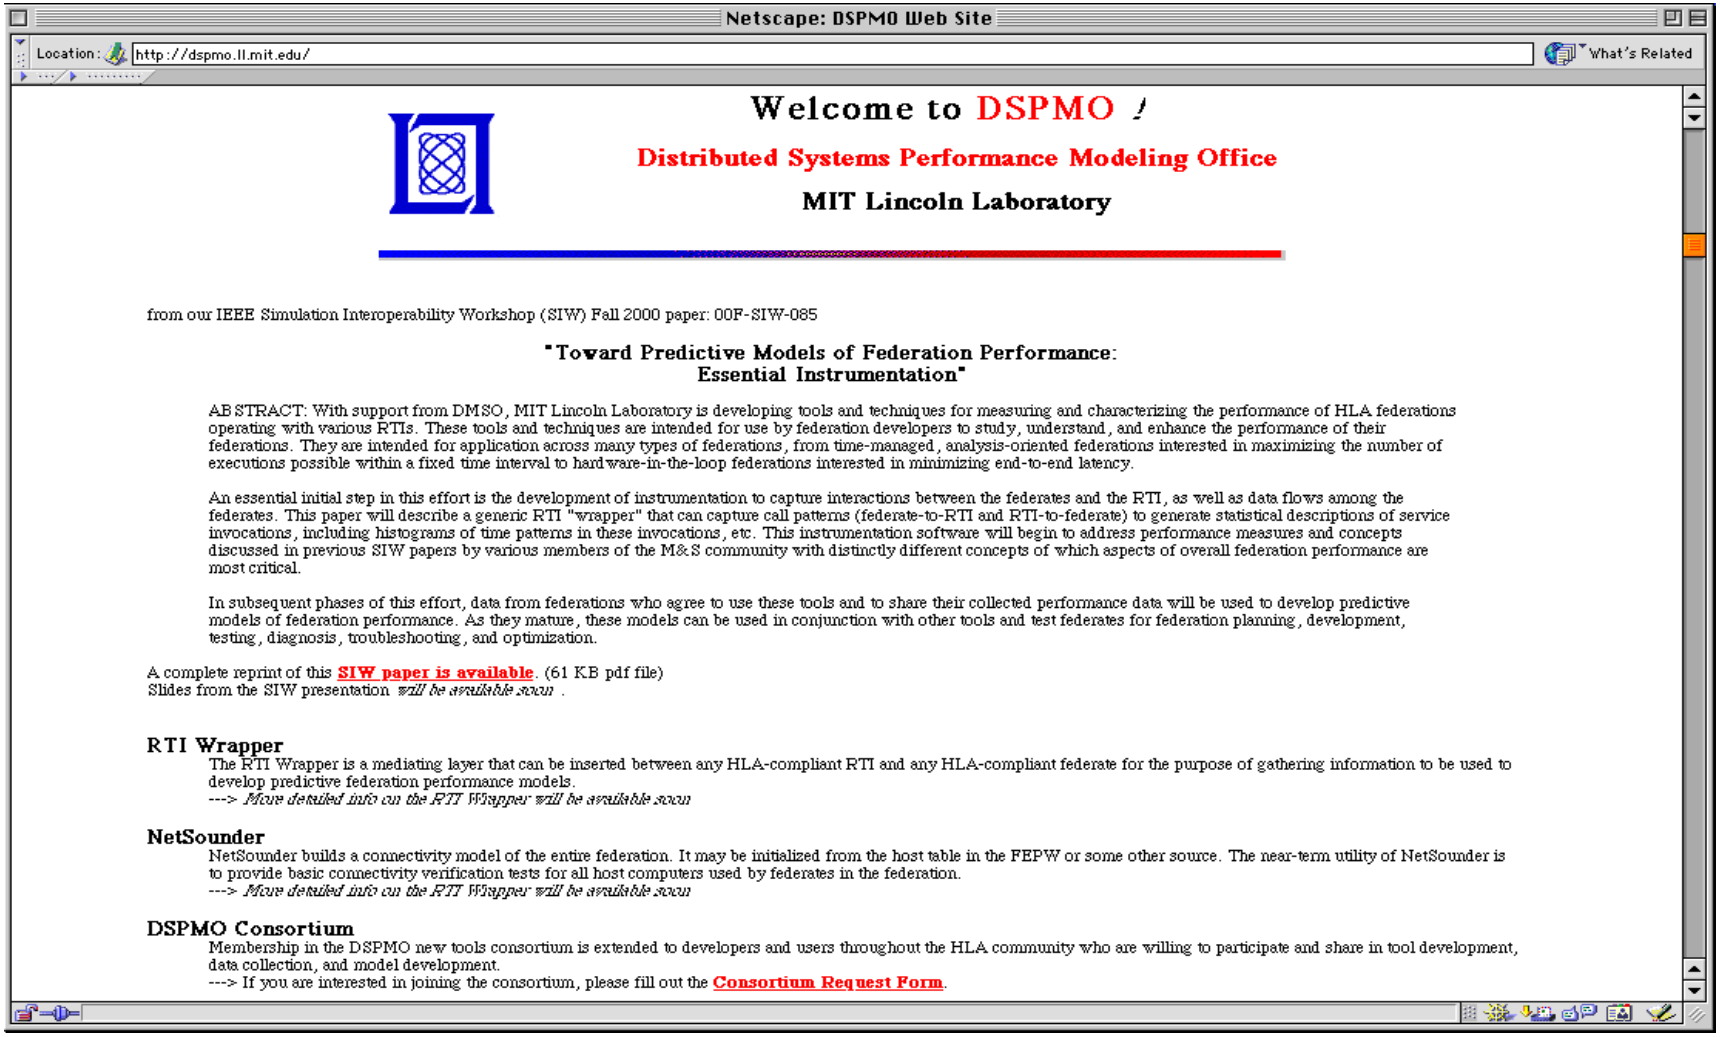

Federation Performance Modeling Consortium

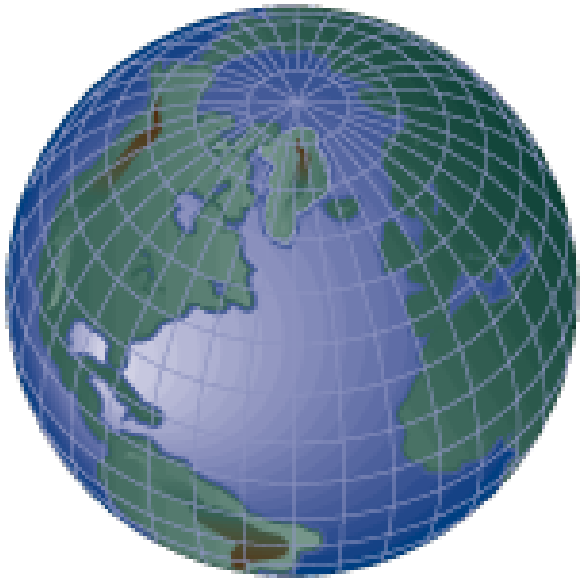

Consortium Operating Principles

- Members agree to share data and models with the Consortium and report the results of their federation performance modeling efforts at SIWs and other appropriate forums
- Members may port Consortium-developed tools to any computing platform/OS combination
  - Any such ports shall be placed in a Consortium repository (with documentation) for use by other members on an "as-is" basis
- Members may extend Consortium-developed tools to support specific applications
  - If such extensions are likely to be useful to other members, developers are encouraged to document them and place them in the repository for use on an "as-is" basis
  - The Consortium cannot promise that any such extensions will be supported in future releases of Consortium tools
- Membership is subject to DMSO approval

Iterative Approach to Federation Performance Modeling

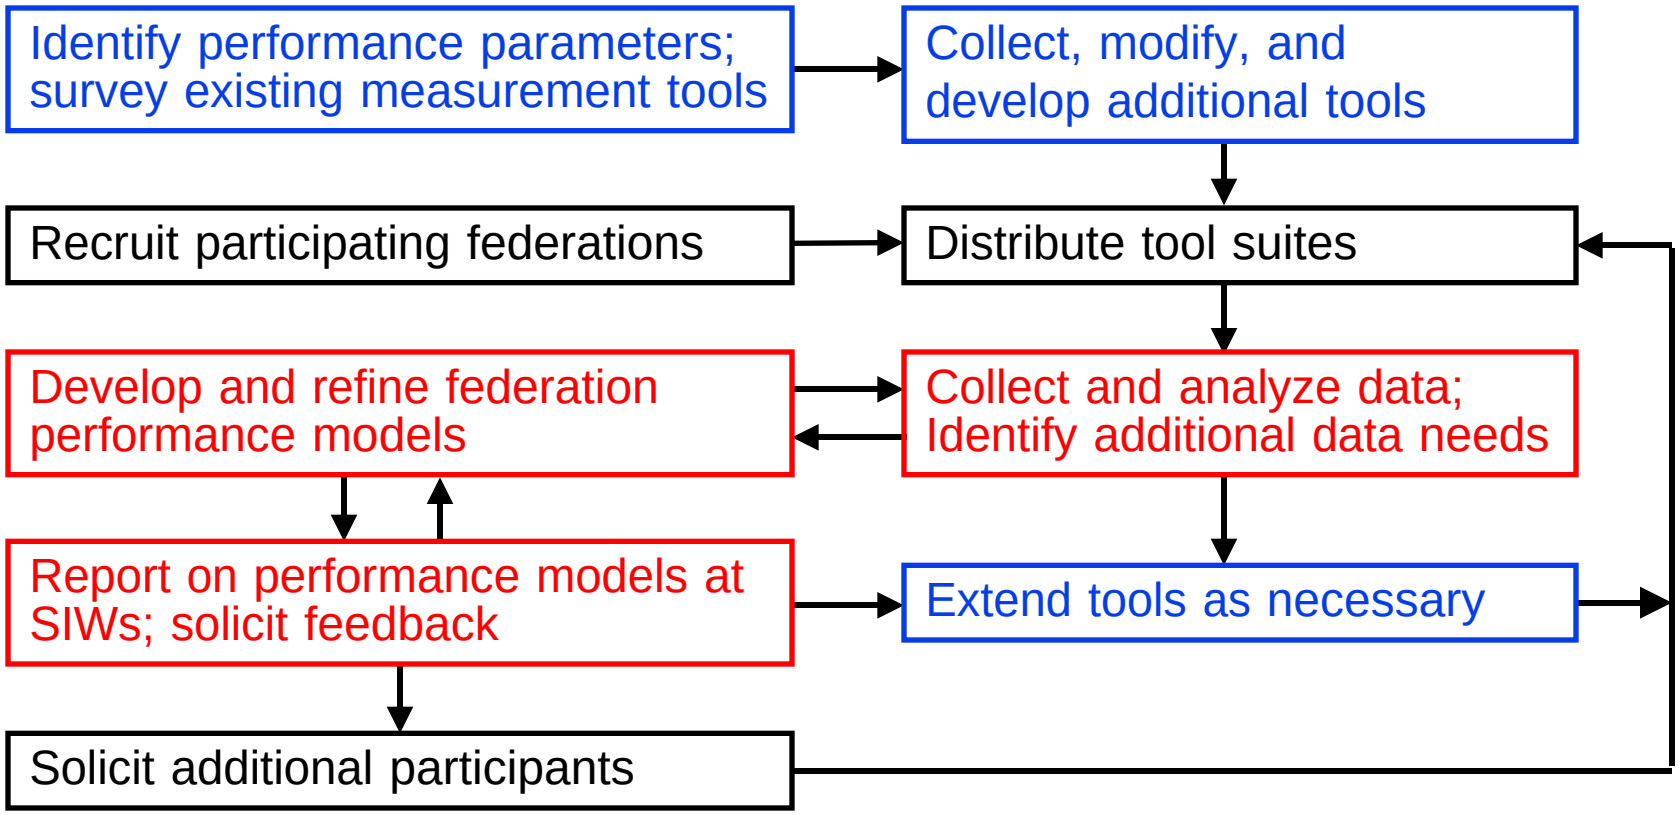

Tool and Model Development Schedule

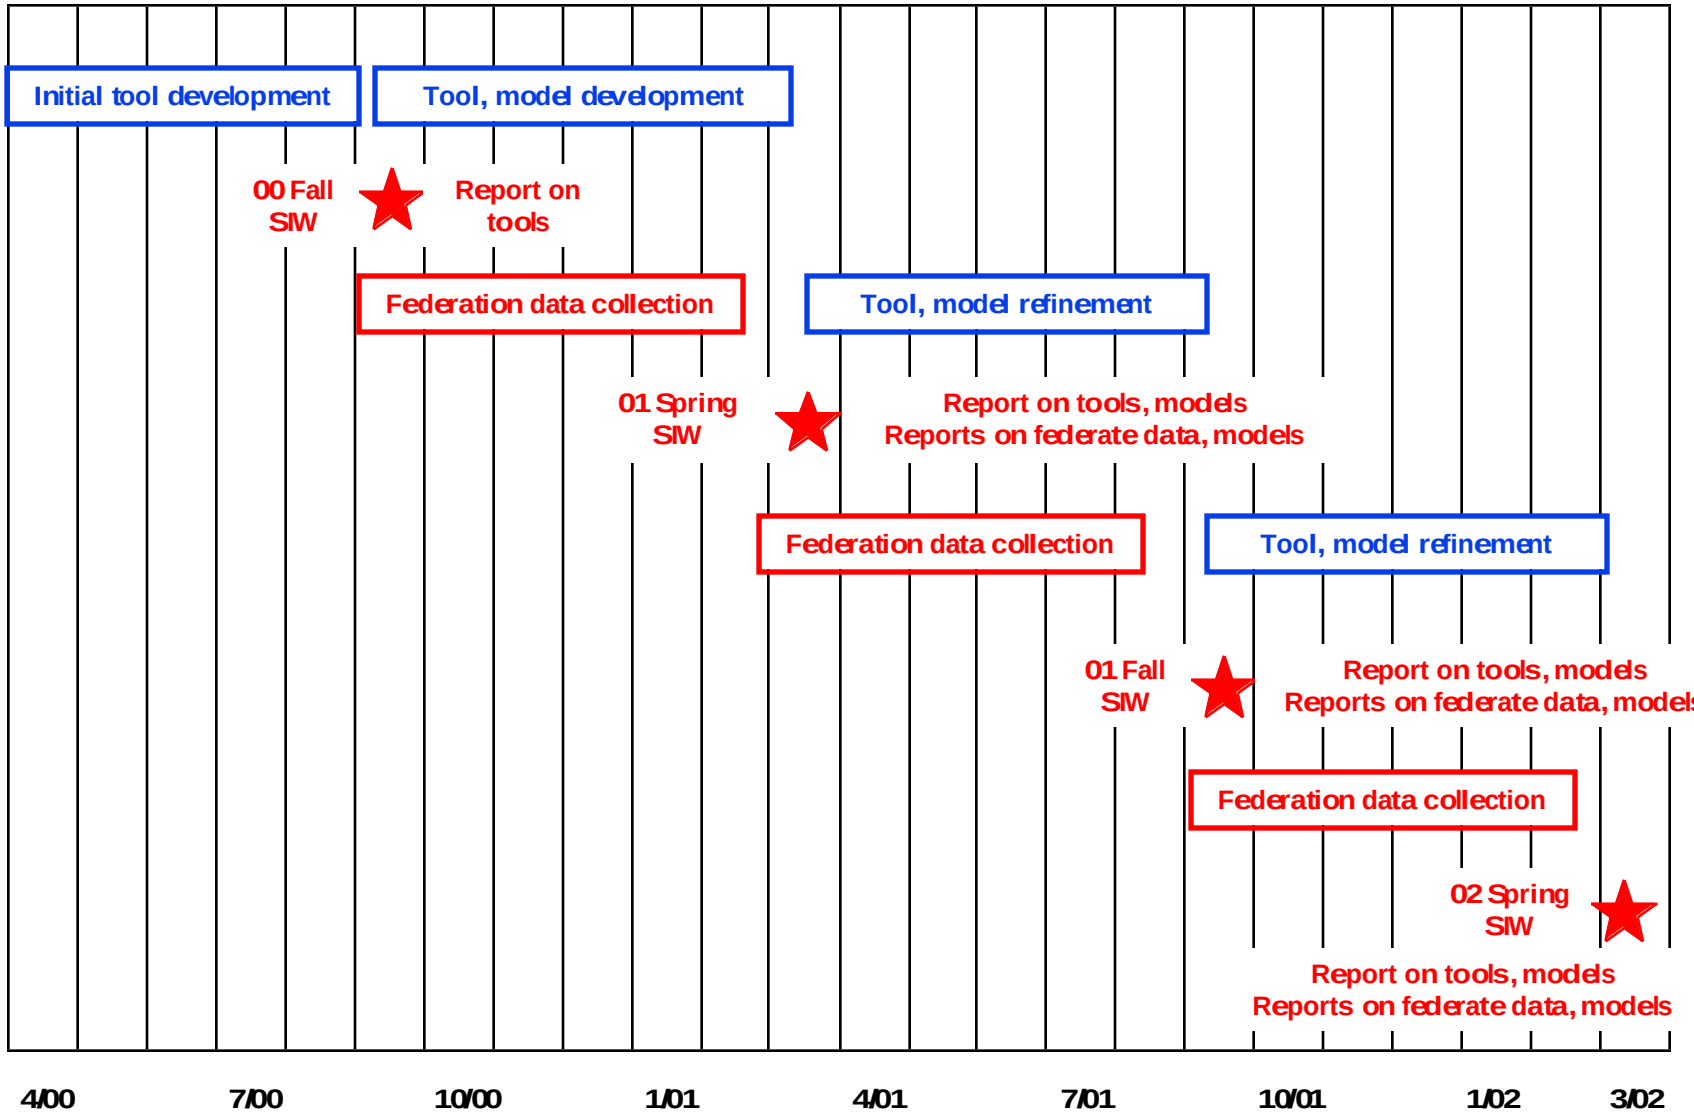

Consortium Web Site

http://DSPMO.LL.MIT.EDU **info@DSPMO.LL.MIT.EDU**