Applications and Implementation of a Mesoscale Numerical Weather Prediction and Visualization System

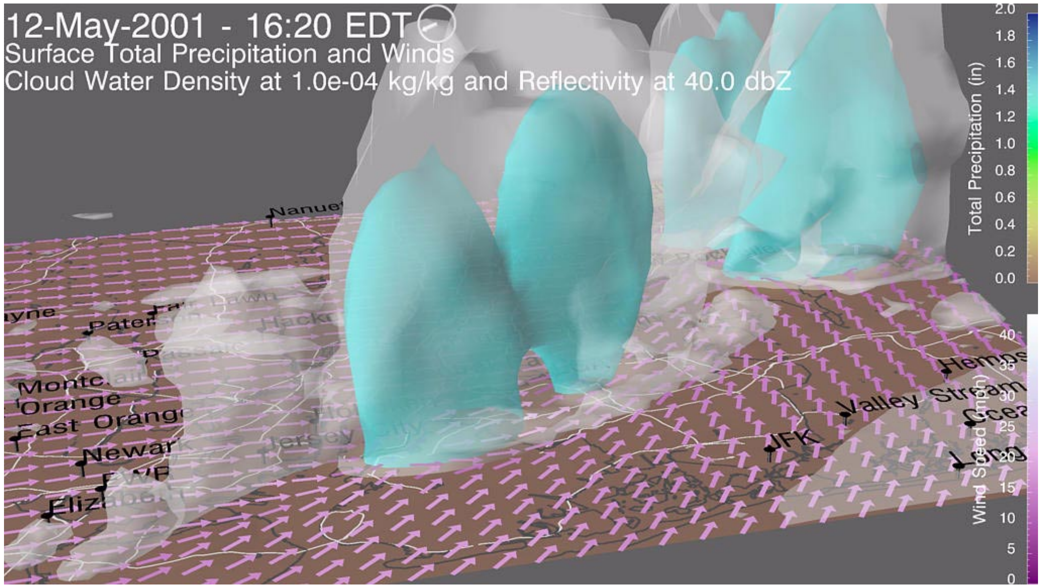

Lloyd A. Treinish and Anthony P. Praino
Mathematical Sciences
IBM Thomas J. Watson Research Center
Yorktown Heights, NY, USA
{lloydt, apraino}@us.ibm.com

http://www.research.ibm.com/weather/DT.html

Applications and Implementation of a Mesoscale Numerical Weather Prediction and Visualization System

- Background and motivation
- Architecture and implementation
- Visualized case study forecasts
  - Interesting severe weather
- Web-based dissemination (current forecast)
- Other examples

Background and Motivation

- Weather-sensitive business operations are often reactive to short-term (3 to 36 hours), local conditions (city, county, state) due to unavailability of appropriate predicted data at this scale
  - Transportation, agriculture, broadcasting, energy, insurance, sports, entertainment, tourism, construction, communications, emergency planning and security warnings
- Mesoscale (cloud-scale) NWP has shown "promise" for years as a potential enabler of proactive decision making for both economic and societal value
- However, . . .

Background and Motivation

- **Despite the "promise" of cloud-scale NWP**
  - Can business and meteorological value be demonstrated beyond physical realism?
  - Can a practical and usable system be implemented at reasonable cost?
- **Improved feasibility today compared to a few years ago due to**
  - Affordable operational computing and visualization platforms
  - Reliability and flexibility of forecasting codes
  - Availability of relevant input data
- **Evaluate concept via a testbed implementation for New York City -- *Deep Thunder***
  - Operational end-to-end infrastructure and automation with focus on HPC, visualization and system integration
  - Forecasts to 1 km resolution for metropolitan area with 3 to 18 hours lead time
  - Prototype business applications with actual end users

New York City Pilot

- Triply nested (62 x 62 x 31) at 16, 4, 1 km (48, 12, 3 sec.)
- Modelling code derived from highly modified version of non-hydrostatic RAMS
- Explicit, full cloud microphysics
- Typically, two 24-hour runs per day (0 and 12 UTC)
- Eta-212/215 via NOAAport for lateral boundaries nudged every 3 hours
- Eta-212/215 for initial conditions after isentropic analysis

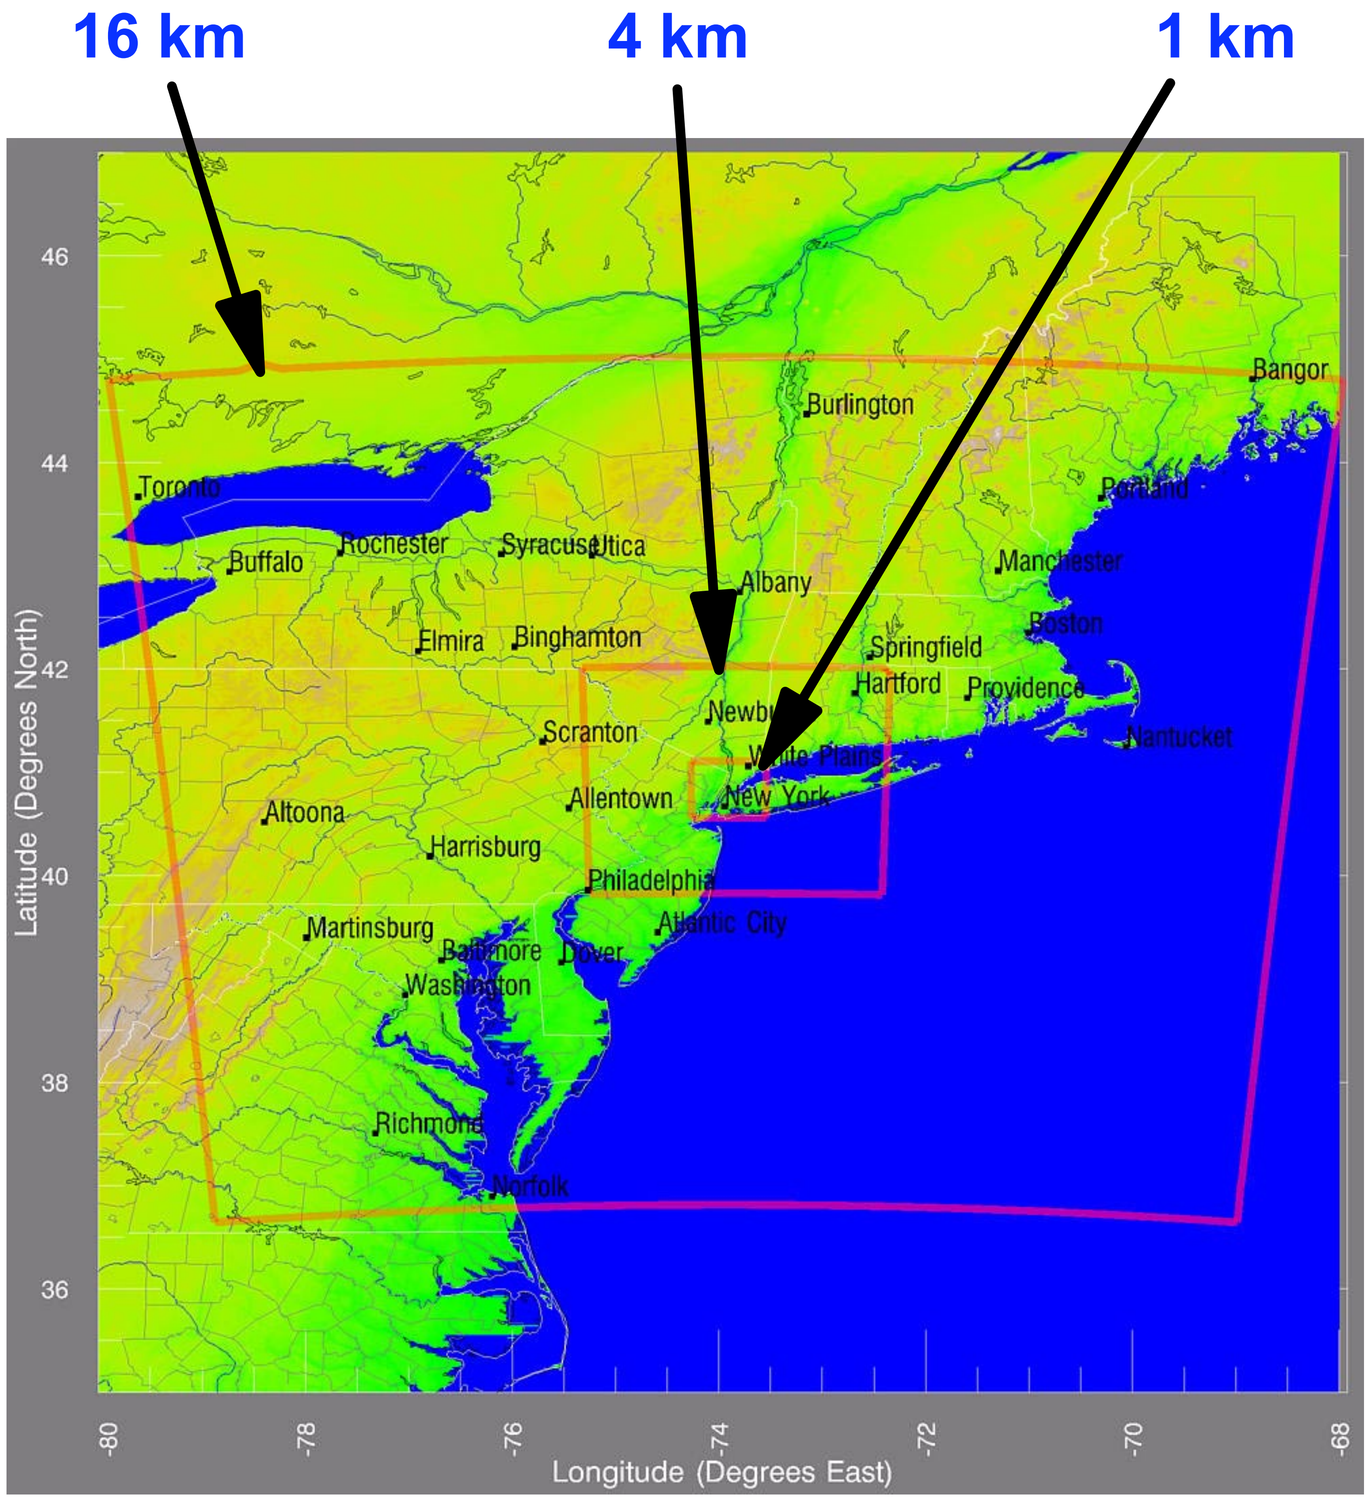

Implementation and Architecture

- Sufficiently fast (>10x real-time), robust, reliable and affordable
  - Focus on HPC, visualization, system integration and automation
- Ability to provide usable products in a timely manner
- Infrastructure needed to be built to enable implementation

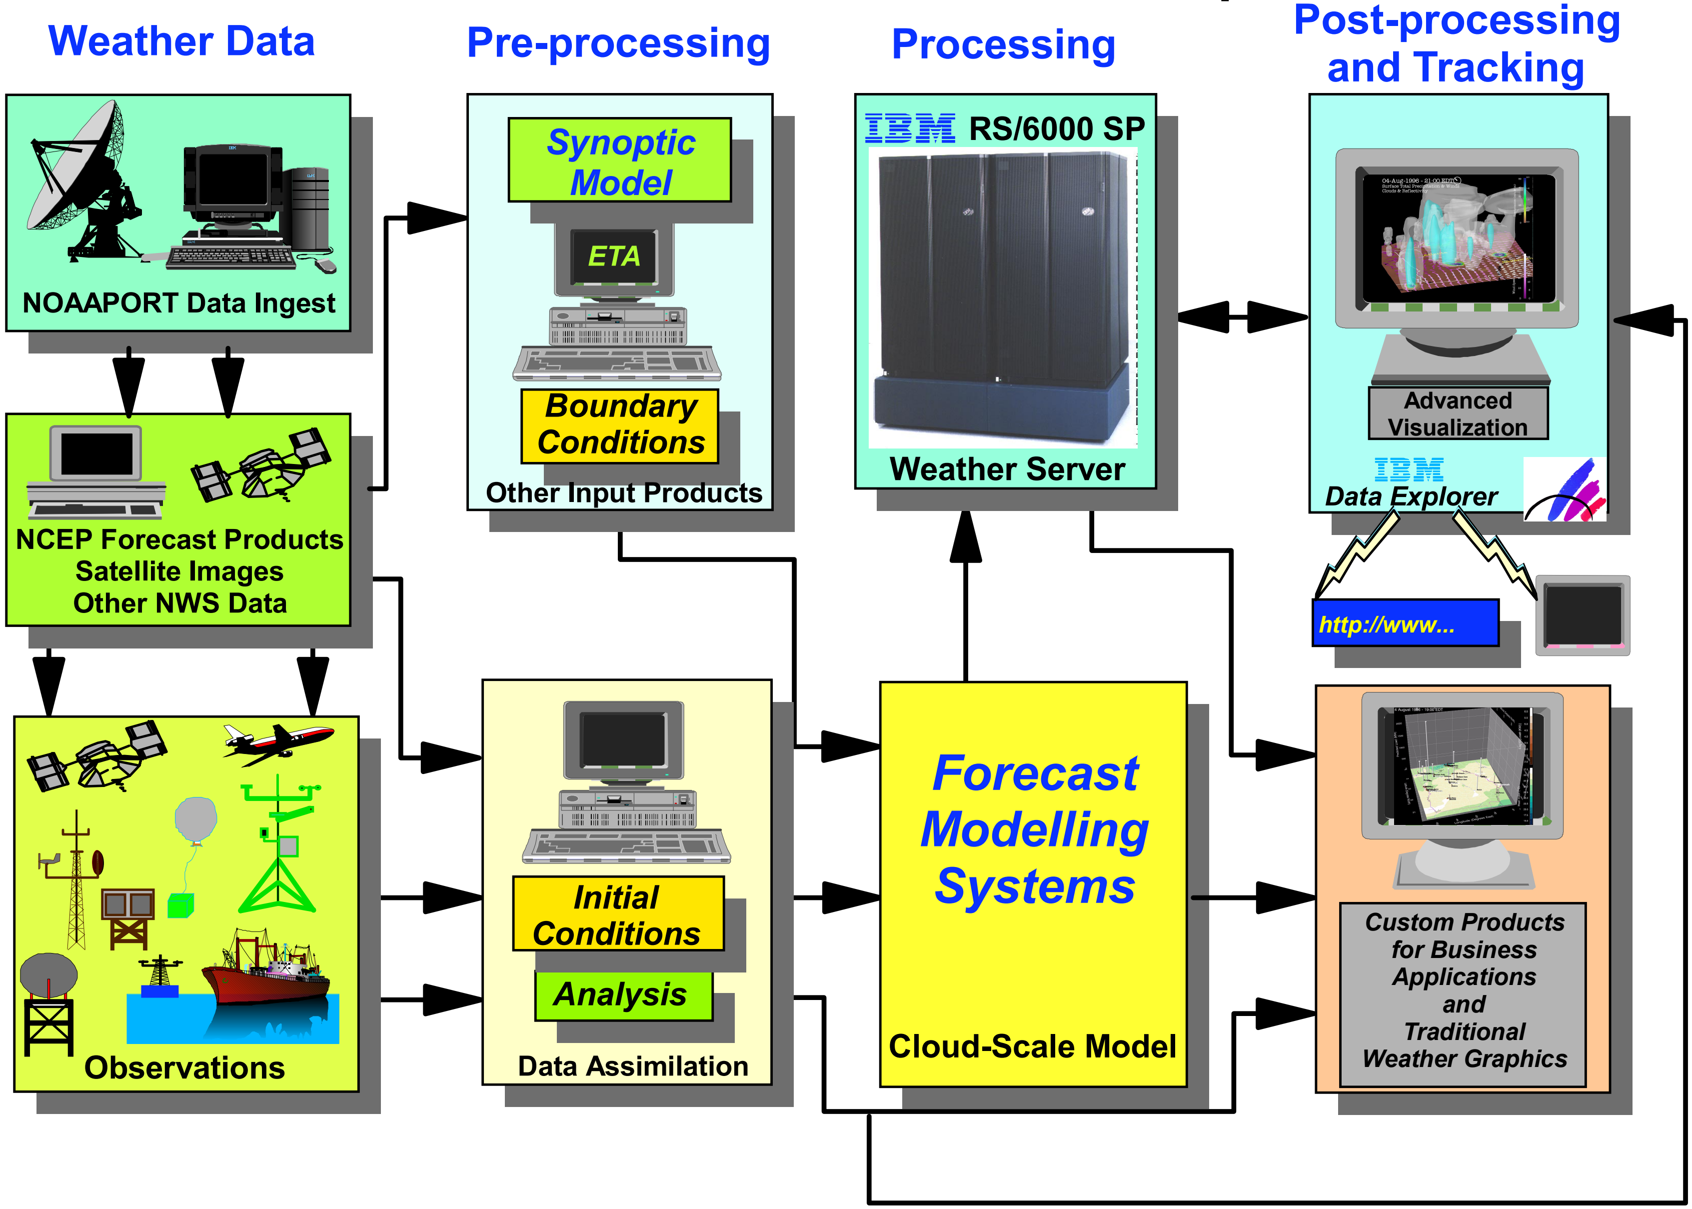

Hardware Environment

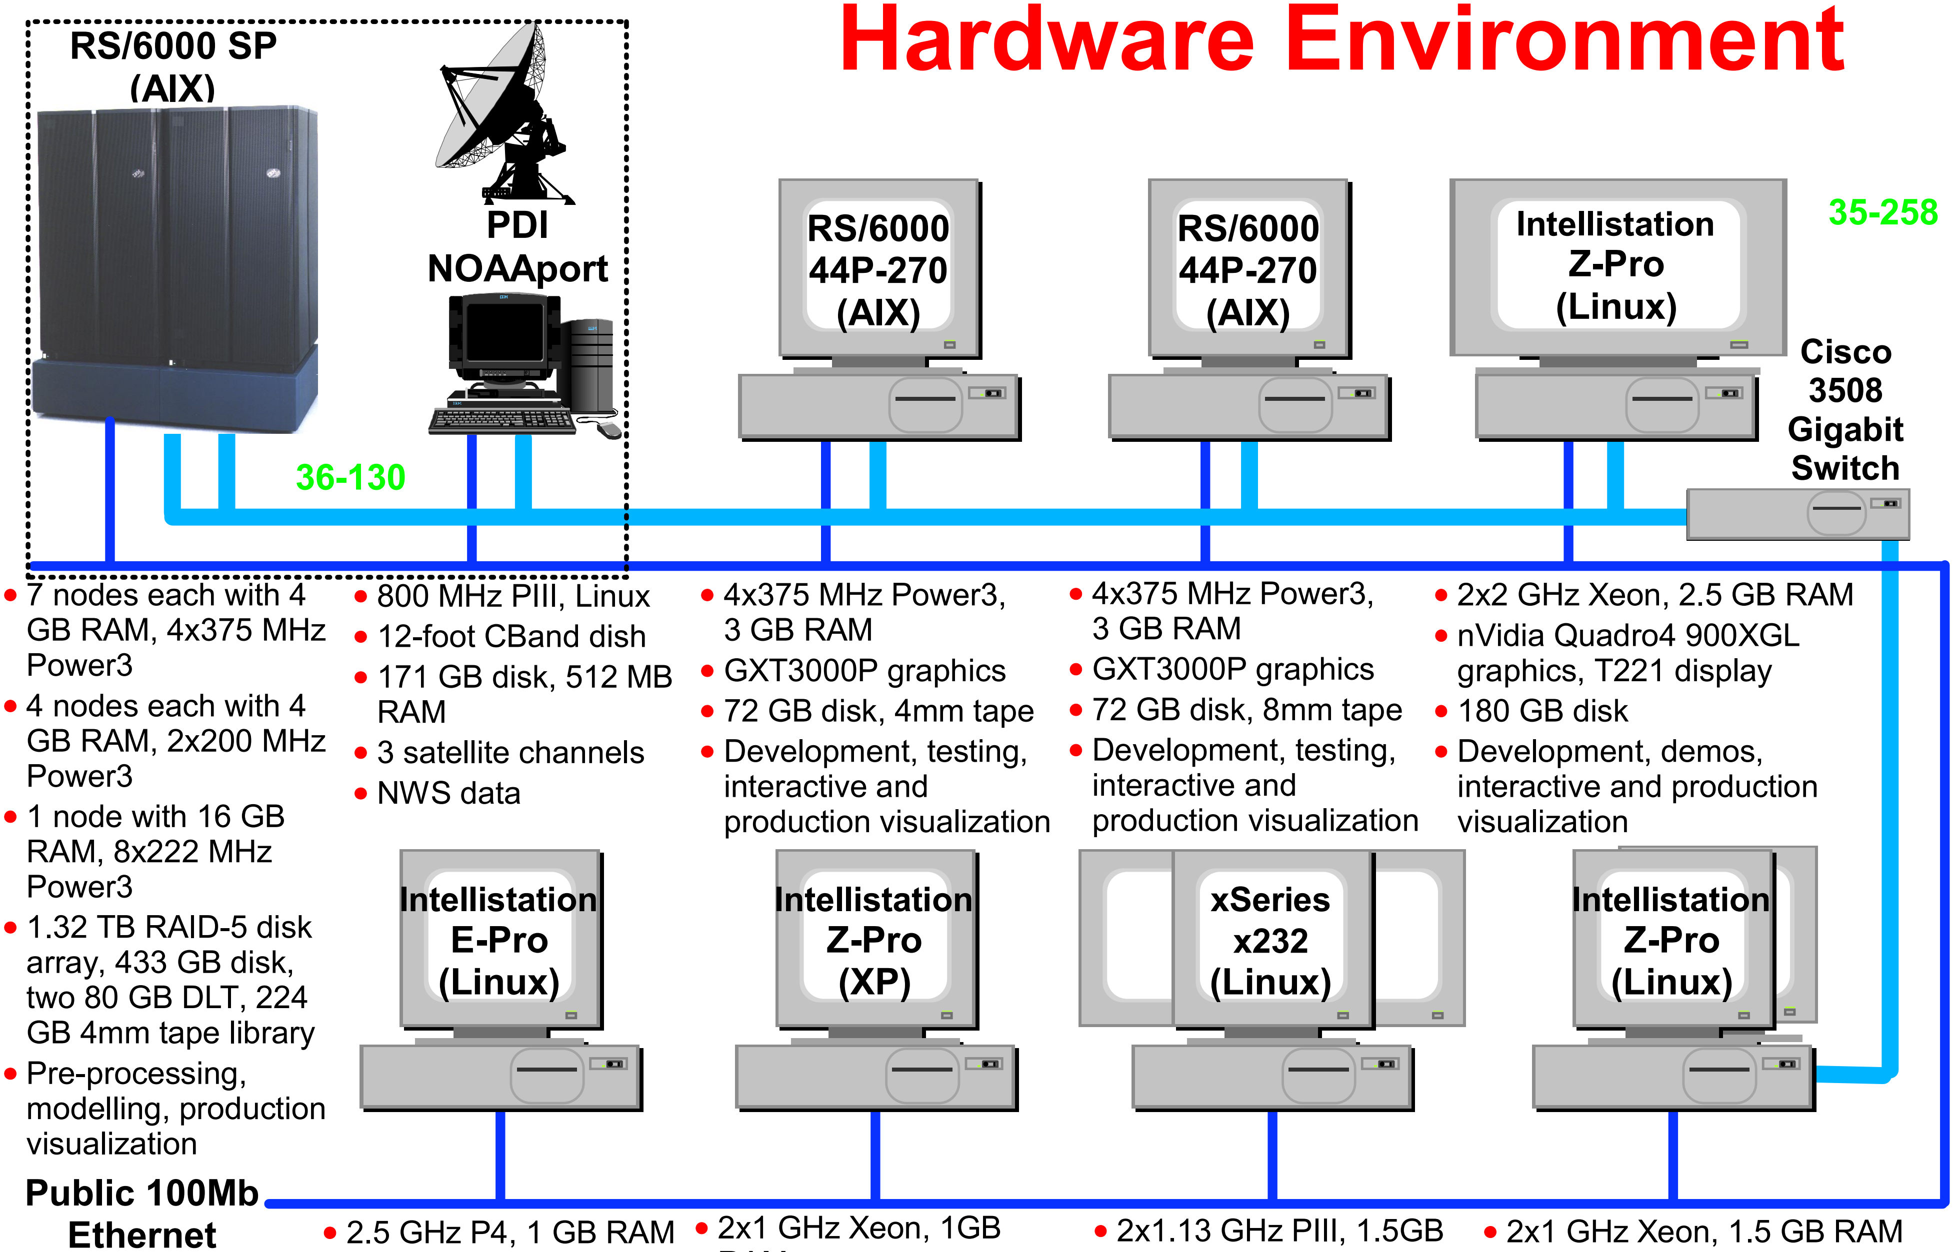

- 7 nodes each with 4 GB RAM, 4x375 MHz Power3
- 4 nodes each with 4 GB RAM, 2x200 MHz Power3
- 1 node with 16 GB RAM, 8x222 MHz Power3
- 1.32 TB RAID-5 disk array, 433 GB disk, two 80 GB DLT, 224 GB 4mm tape library
- Pre-processing, modelling, production visualization

- 800 MHz PIII, Linux
- 12-foot CBand dish
- 171 GB disk, 512 MB RAM
- 3 satellite channels
- NWS data

- 4x375 MHz Power3, 3 GB RAM
- GXT3000P graphics
- 72 GB disk, 4mm tape
- Development, testing, interactive and production visualization

- 4x375 MHz Power3, 3 GB RAM
- GXT3000P graphics
- 72 GB disk, 8mm tape
- Development, testing, interactive and production visualization

- 2x2 GHz Xeon, 2.5 GB RAM
- nVidia Quadro4 900XGL graphics, T221 display
- 180 GB disk
- Development, demos, interactive and production visualization

- 2.5 GHz P4, 1 GB RAM
- nVidia Quadro4 580XGL graphics
- 120 GB disk
- Demos

- 2x1 GHz Xeon, 1GB RAM
- FireGL4 graphics
- 117 GB disk, (2) CD-RW, Jaz, Zip, scanner
- Demos, interactive visualization, development

- 2x1.13 GHz PIII, 1.5GB RAM
- 3 nVidia GeForce2 MX graphics triple display
- 72 GB disk
- Development, interactive and production visualization

- 2x1 GHz Xeon, 1.5 GB RAM
- FireGL4 graphics, dual monitors
- 81 GB disk, 8mm tape
- Development, NWS graphics, interactive and production visualization

Forecast Pre-Processing Procedural and Data Flow

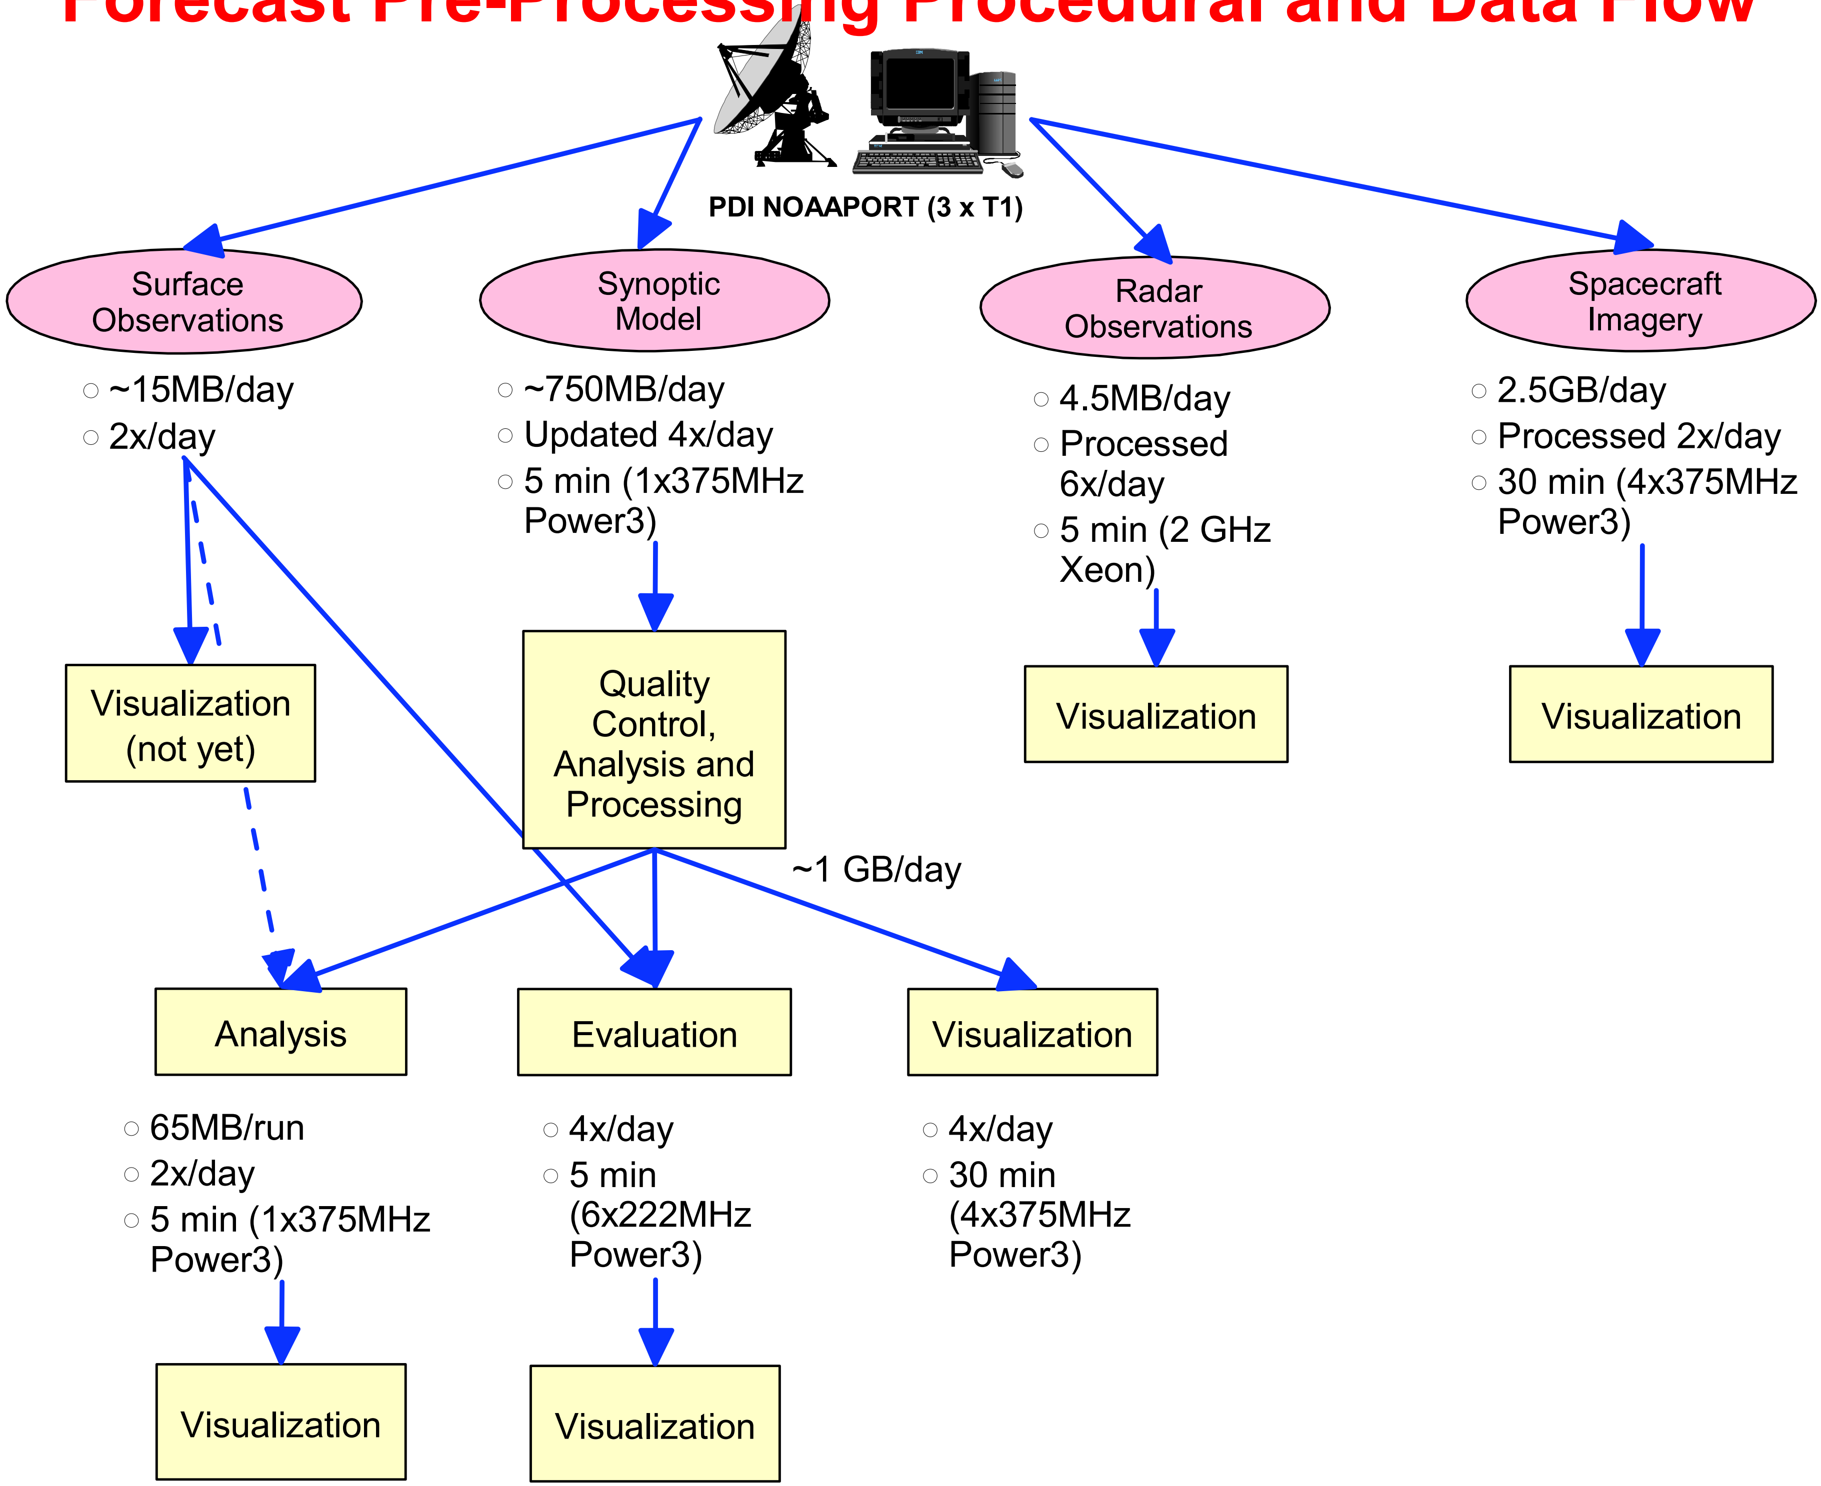

Visualization Component

- Traditional meteorological visualization is typically driven by data for analysis -- inappropriate for business applications
- Timely usability of cloud-scale NWP results requires
  - Understanding of how weather data are used
  - Identification of user goals, which are mapped to visualization tasks
  - Mapping of data to visualization tasks
  - Forecaster has control over content by design or simple interaction
  - Non-forecaster has limited control over content (targeted design) and simple interaction
  - Products designed in terms relevant for user
- Approach and implementation as presented at the ECMWF 2001 workshop is integrated operationally
- Wide range of generic capabilities needed
  - Line plots to 2d maps to 3d animations -- but customized
  - Assessment, decision support, analysis and communications
  - Automated (parallelized) generation of products for web dissemination
  - Highly interactive applications on workstations

Forecast Post-Processing Procedural & Data Flow

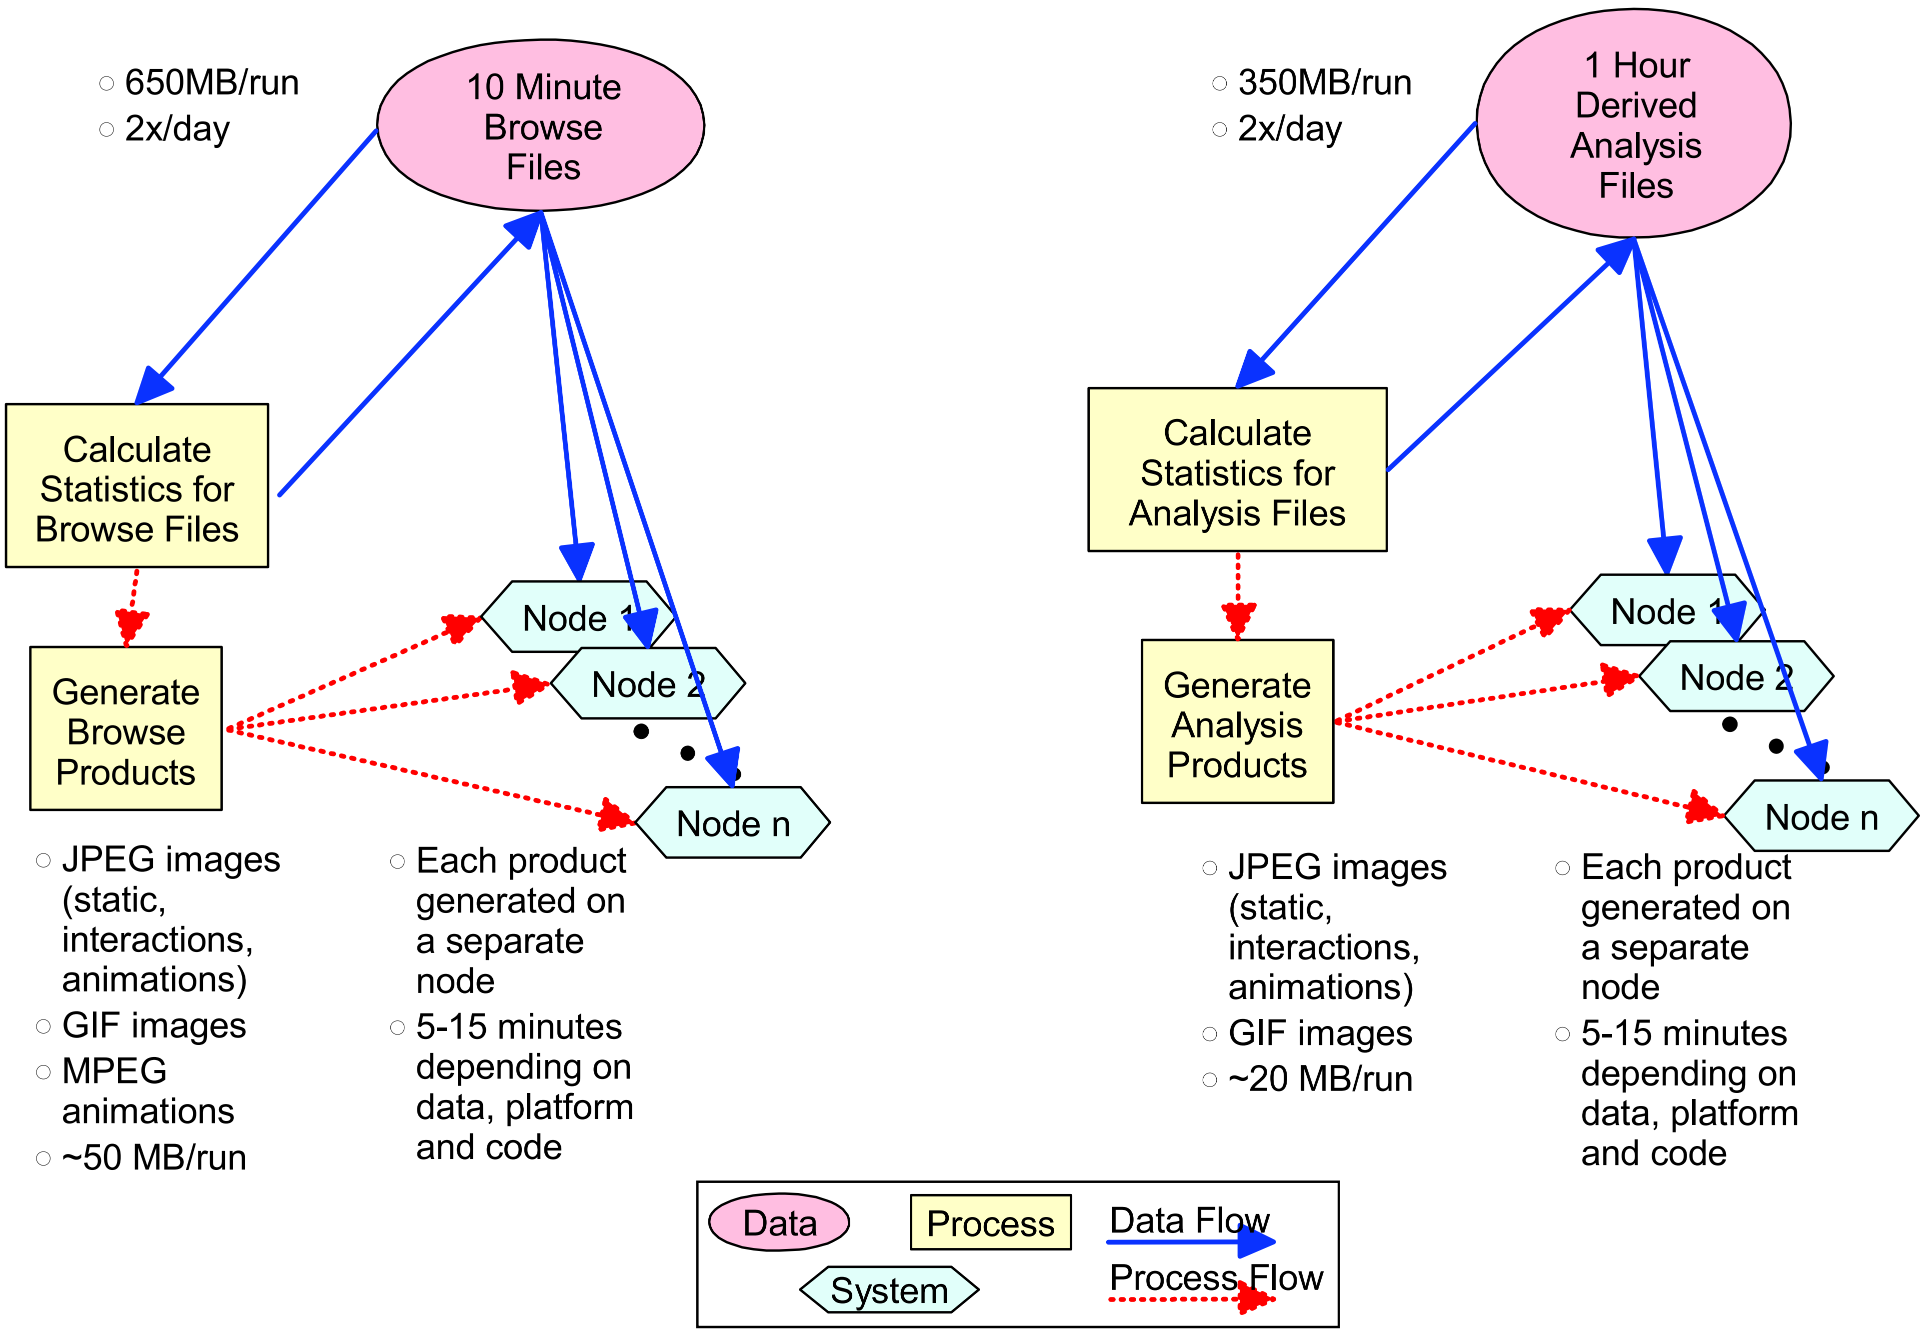

Visualized Case Studies

- **Severe convective events**
- **Unusual spring snowfall**
- **Heavy snowstorm**

Severe Convective Event
31 May 2002 -- Early Evening

- **In the late afternoon and early evening, a strong cold front clashed with a very unstable air mass across eastern New York State coupled with strong and divergent winds aloft**
- **Resulted in widespread severe damage to buildings plus downed trees and power lines in eastern New York State -- mostly strong winds associated with thunderstorms as well as large hail**
- **Funnel cloud visible in 1km nest while two F1 tornados confirmed north of 4km nest (in 16km nest)**

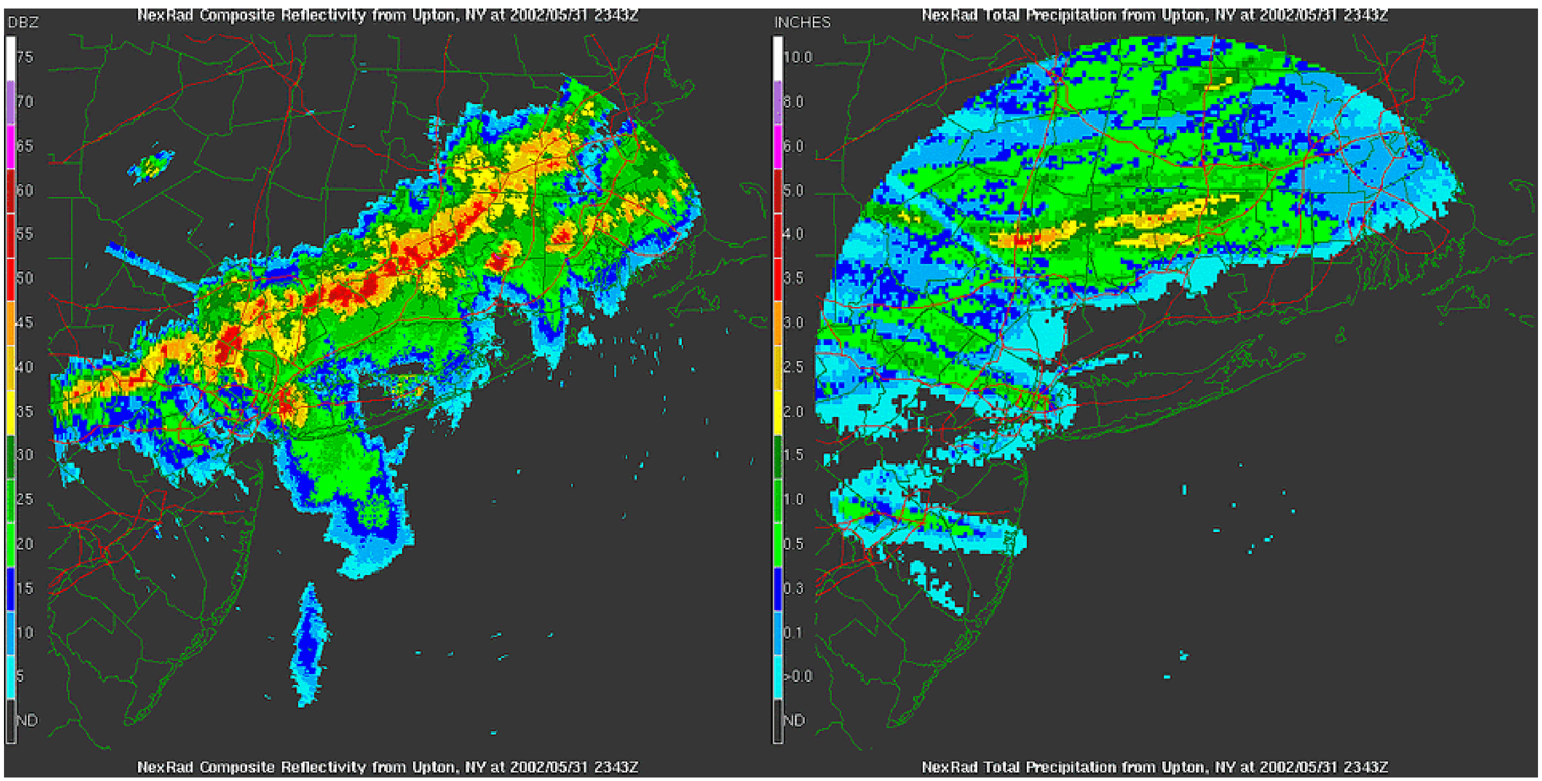

Forecast Results
31 May 2002 -- Early Evening

- **Forecast initiated with data from 12 UTC**
- **Results available about 17 UTC, > 6 hour lead time**
- **Negative bias in timing ~25 minutes**
- **Intensity roughly correct**
- **Missed southeastern portion (beyond 4km domain)**

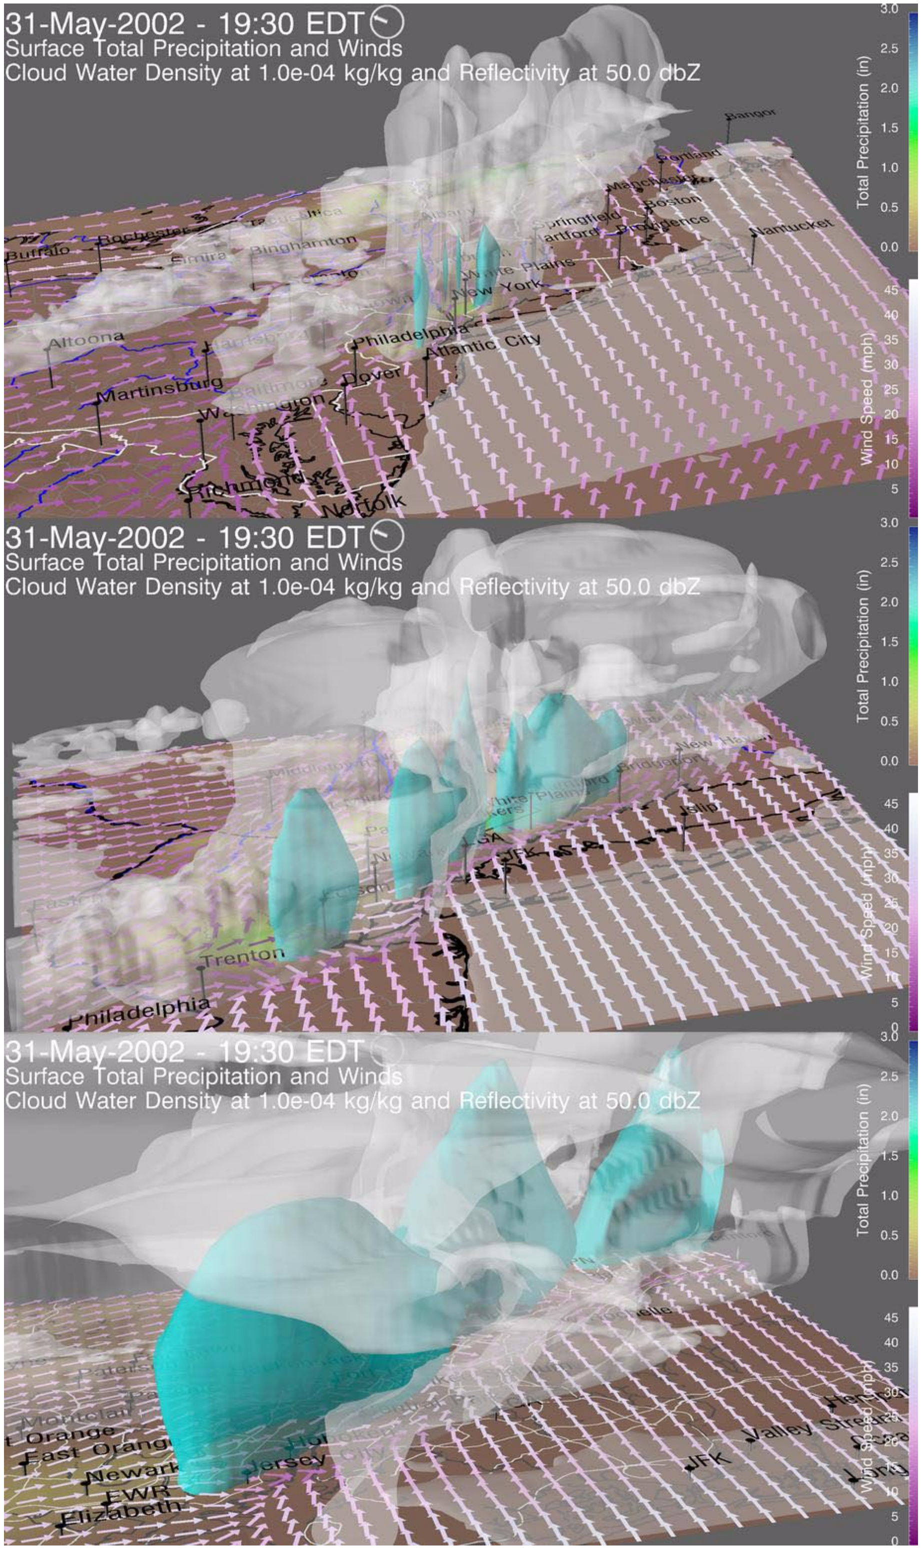

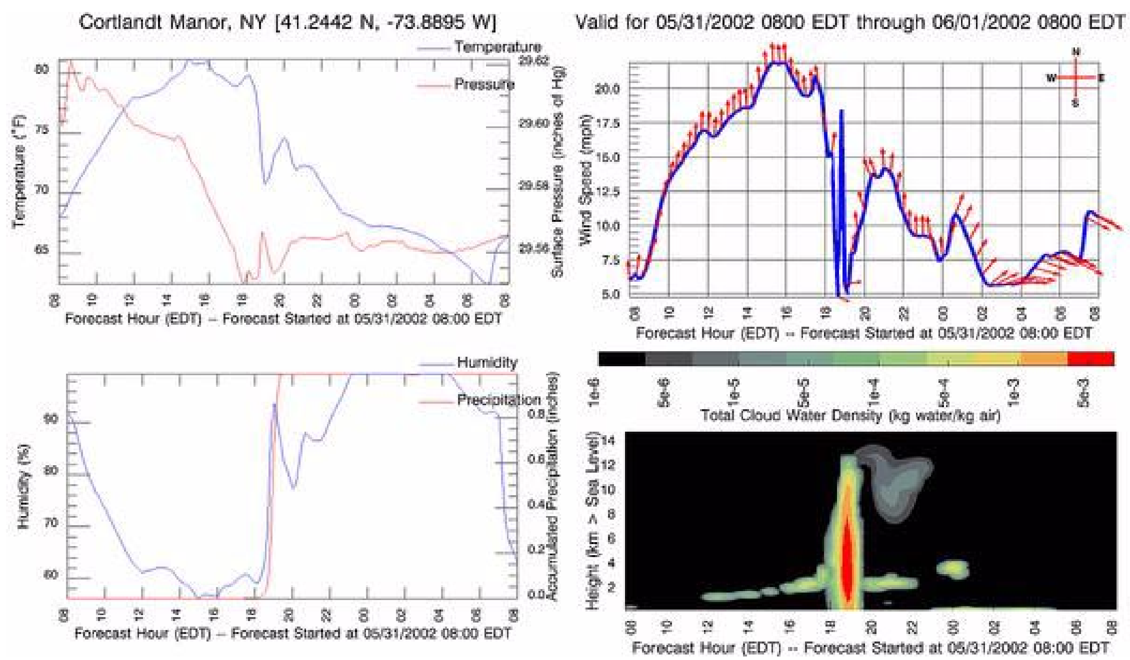

Severe Convective Event
23 September 2003 -- Early Morning

- **Weak cold front in Pennsylvania and New Jersey clashed with a surprisingly unstable air mass**
- **Resulted in 100 trees uprooted, downed power lines and damage to buildings in New Jersey**
- **Gusts up to 130 km/hr and heavy rain (> 60mm in 30 minutes)**
- **Two F1 tornados confirmed in 4km nest (New Jersey) between 1205 and 1235 UTC, two additional F1s confirmed in 16km nest (Pennsylvania)**

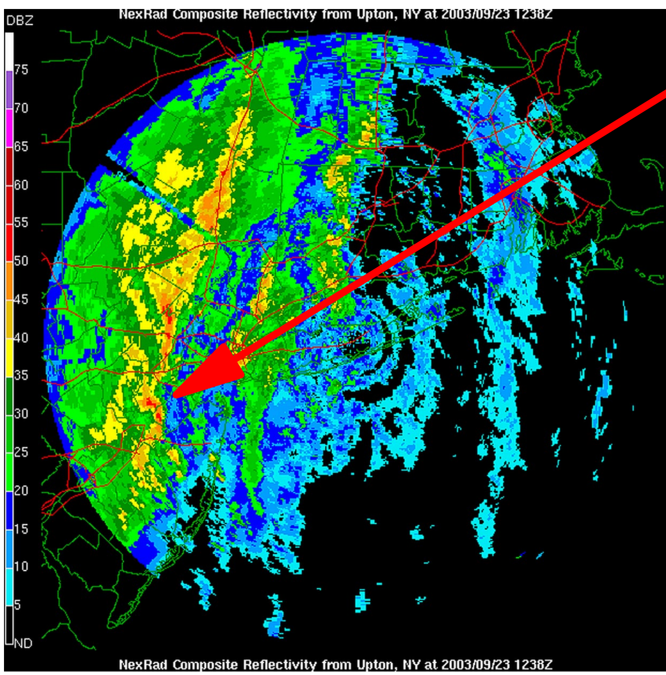

Radar Imagery -- 0838 EDT

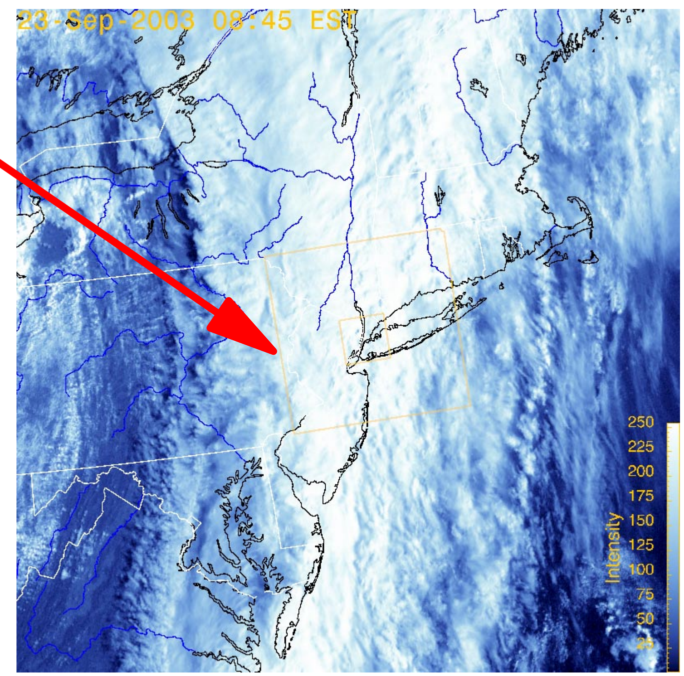

Cloud Imagery -- 0845 EDT

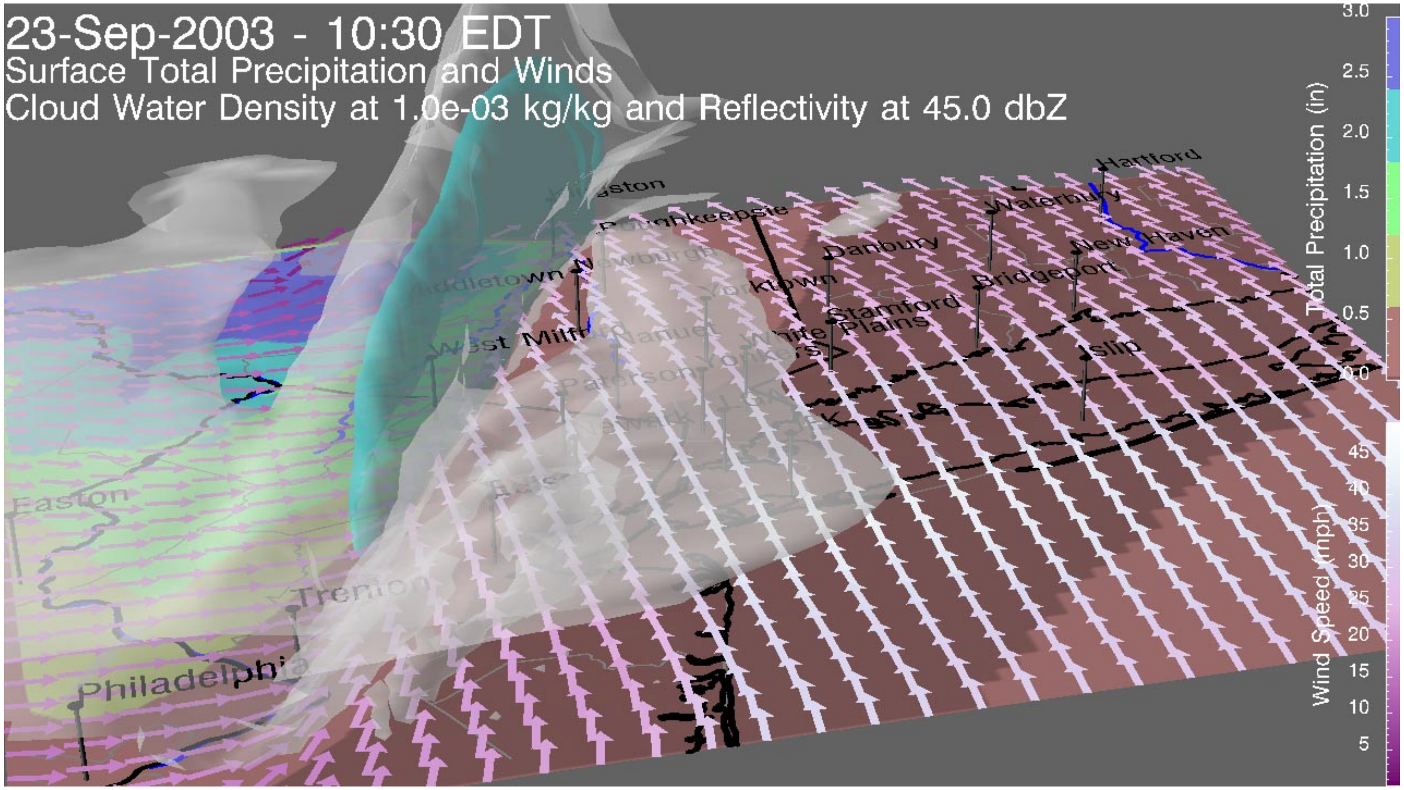

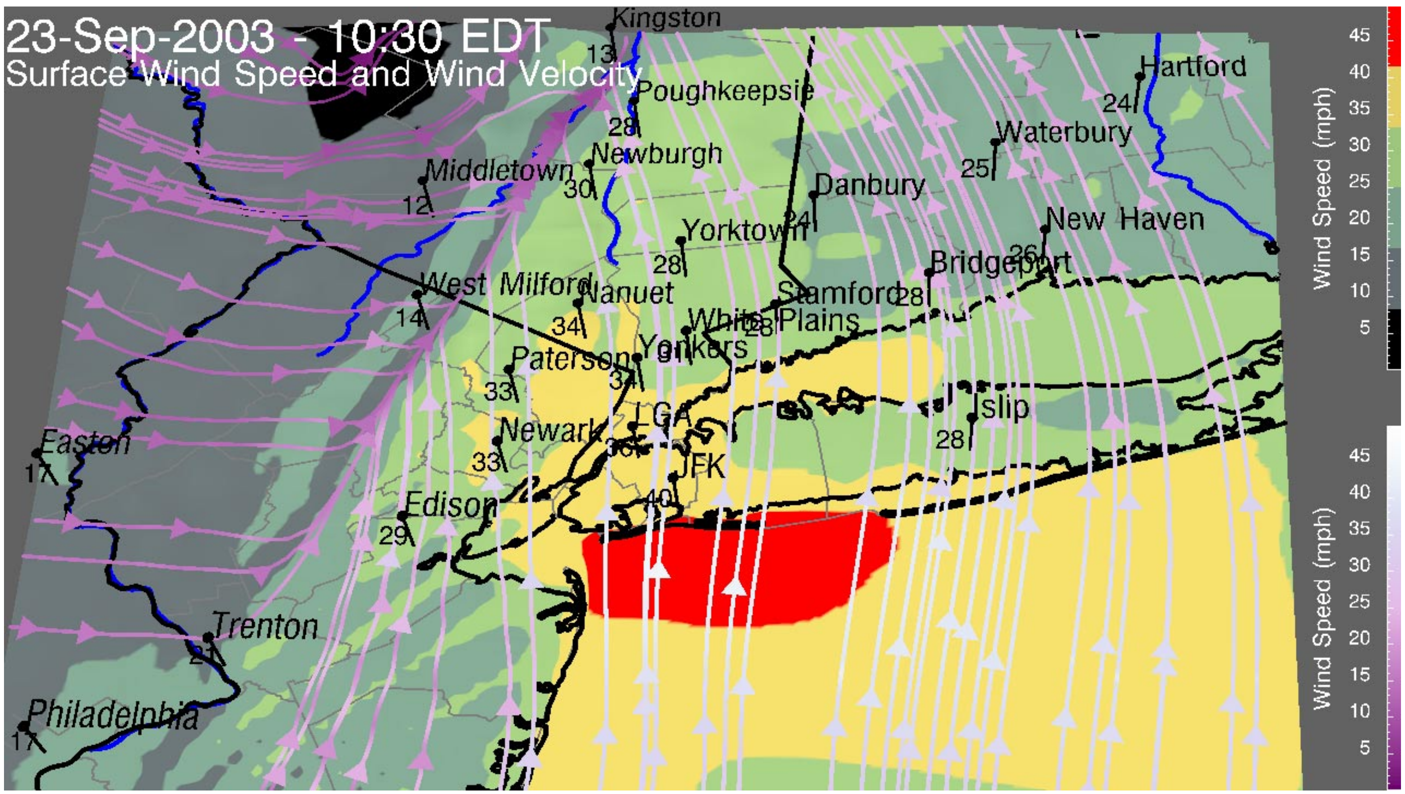

Forecast Results 23 September 2003 Early Morning

- **Well-formed squall line is forecasted, associated with strong convection in northern New Jersey after 12 UTC with the potential for tornado genesis**
- **Location biased to the north and positive bias in timing of 90 to 120 minutes**
- **Despite error, significant "heads-up" for event (> 7 hours) with roughly correct intensity**
- **Forecast initiated with data from 0 UTC with results available about 5 UTC**

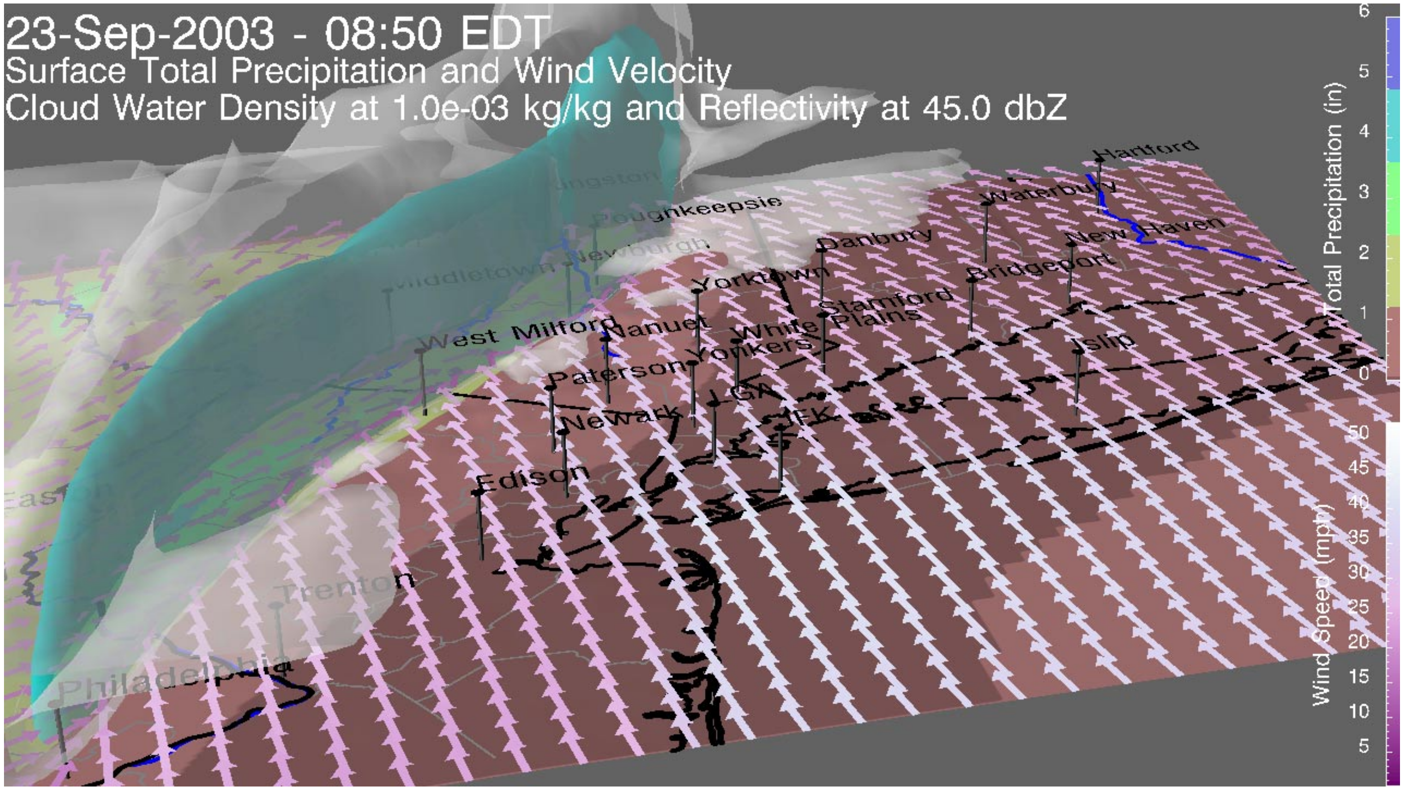

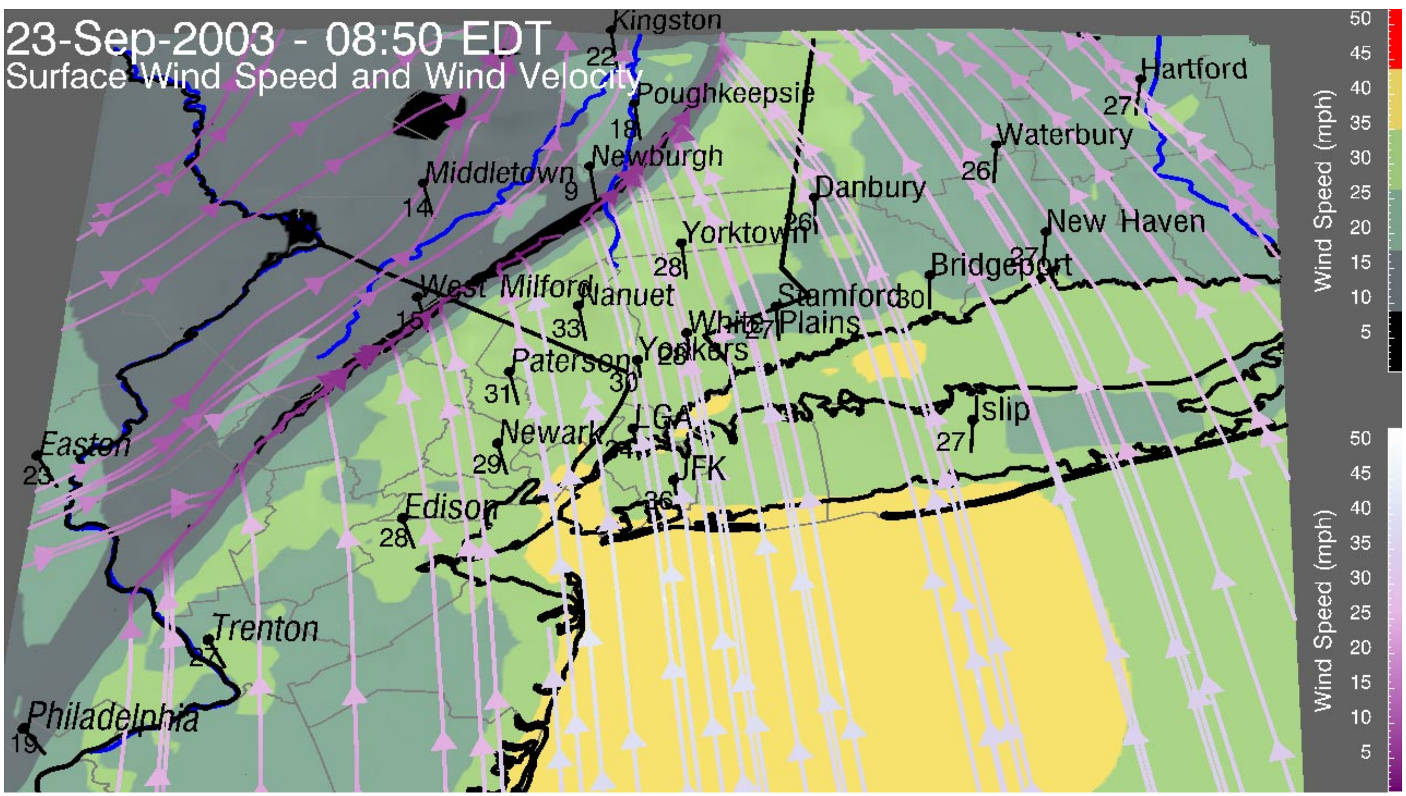

Forecast Results
23 September 2003
Early Morning

- **Could a later forecast have provided better results ?**
- **Hindcast initiated with data from 6 UTC, which operationally would have been available by 1030 UTC**
- **Bias in location and time corrected, which would have still provided a lead time for the event**

Late Spring Snowfall 18 May 2002 -- Early Morning

- An area of low pressure developed over the middle Atlantic states, which tracked northeast, off into the Atlantic Ocean

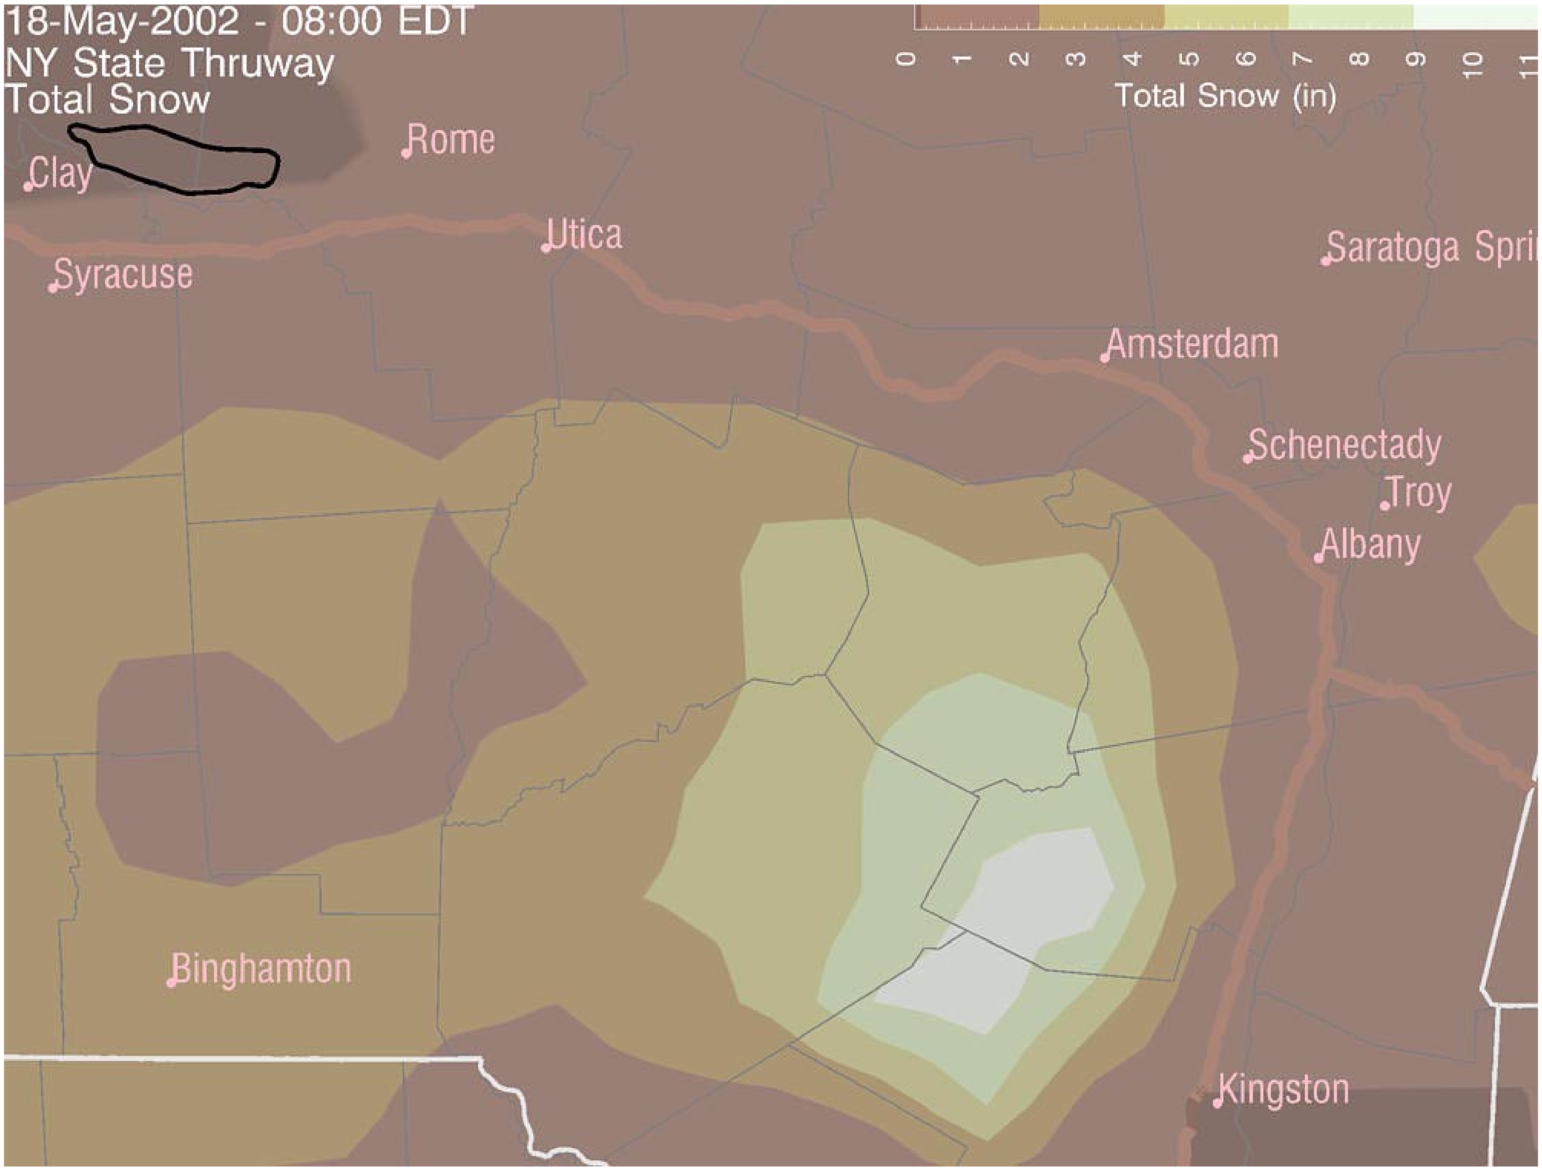

- The air north of the storm was marginally cold enough to turn rain into snow across eastern New York State
- Latest measurable snow in Albany (5.5 cm) with reports up to 20 cm
- Forecast initiated with data from 12 UTC, available about 17 UTC, > 15 hour lead time
- Snow amounts may have positive bias

Blizzard -- February 17, 2003

- An exceptionally strong (Arctic) high pressure system pushing over New England, a low pressure system developed in the Tennessee Valley and tapped into moisture from the Gulf of Mexico, and then redeveloped rapidly off the middle Atlantic coast
- Periods of light snow developed as northeast winds increased to around 15 mph across the New York City area on the afternoon of February 16th
- Snow became widespread and heavy, falling at rates up to 2 to 3 inches per hour that night and the next day
- Heavy snow blown by northeast winds 20 to 30 mph caused near blizzard conditions throughout the area, crippling transportation
- In New York City, the cost estimate for total snowfall operations was ~$20M
- Widespread moderate beach erosion and minor tidal flooding occurred
- The storm tracked northeast and was east of Cape Cod by early February 18
- Storm totals ranged from about 30 to 70 cm

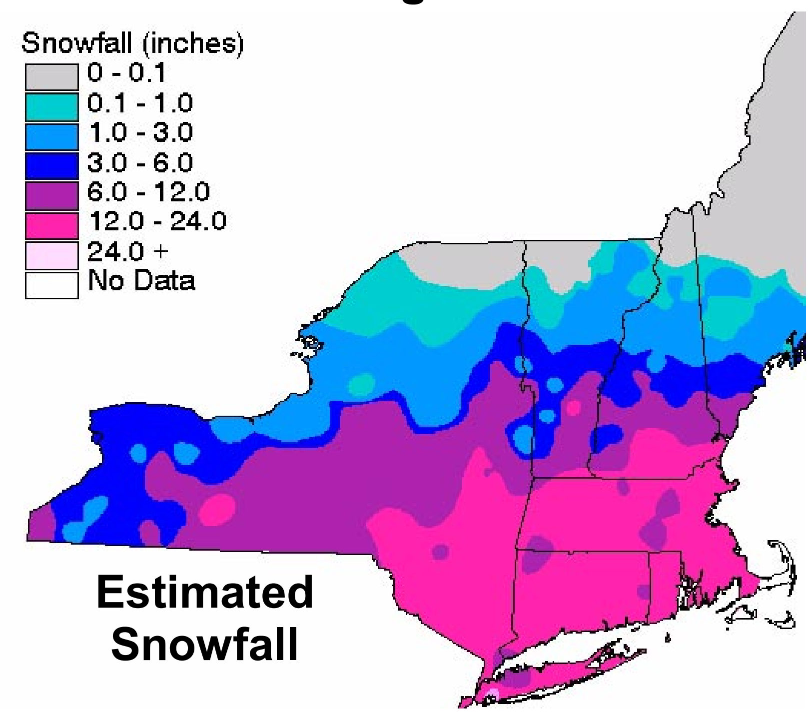

Measured Snowfall (Inches)

JFK	25.6	Middletown	15.0
Nyack	18.0	LGA	16.5
C. Park	19.8	White Plains	17.0
Upton	21.1	Bridgeport	17.0
Yorktown	21.0	W. Milford	28.0
Yonkers	19.0	EWR	22.1
Islip	14.0	Danbury	19.5

Forecast Results
February 17, 2003

- **Forecast initiated with data from 0 UTC**
- **Results available about 5 UTC, hours before heavy snow began**
- **Good agreement in both snow totals ("dry" algorithm) and geographic distribution**

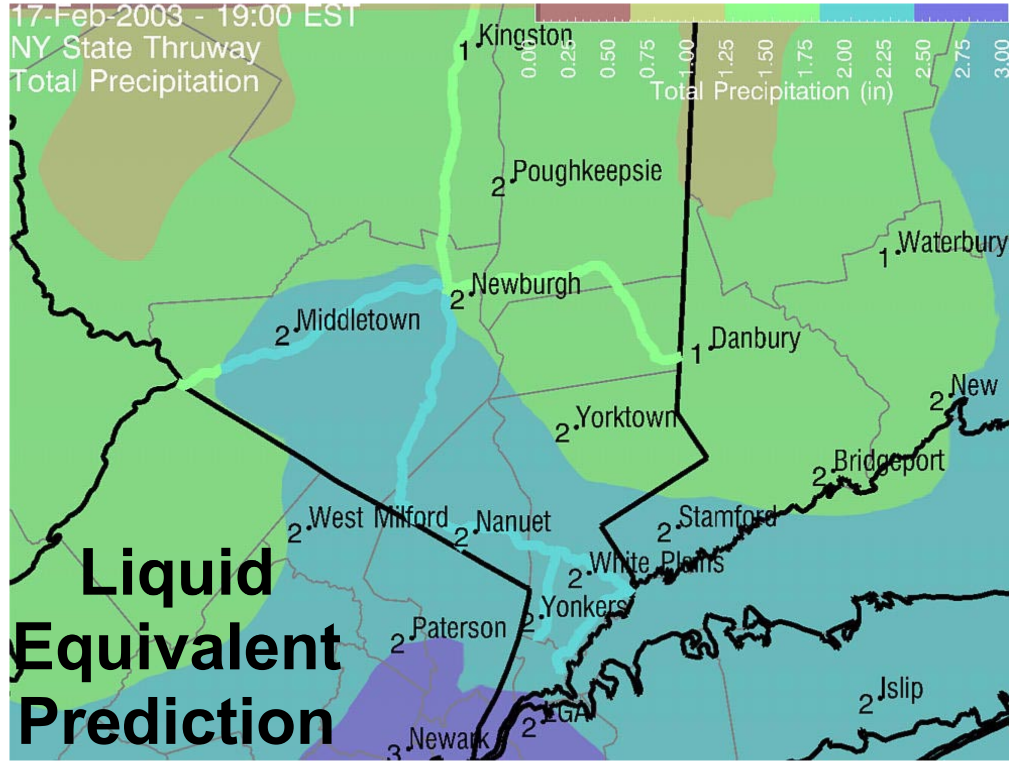

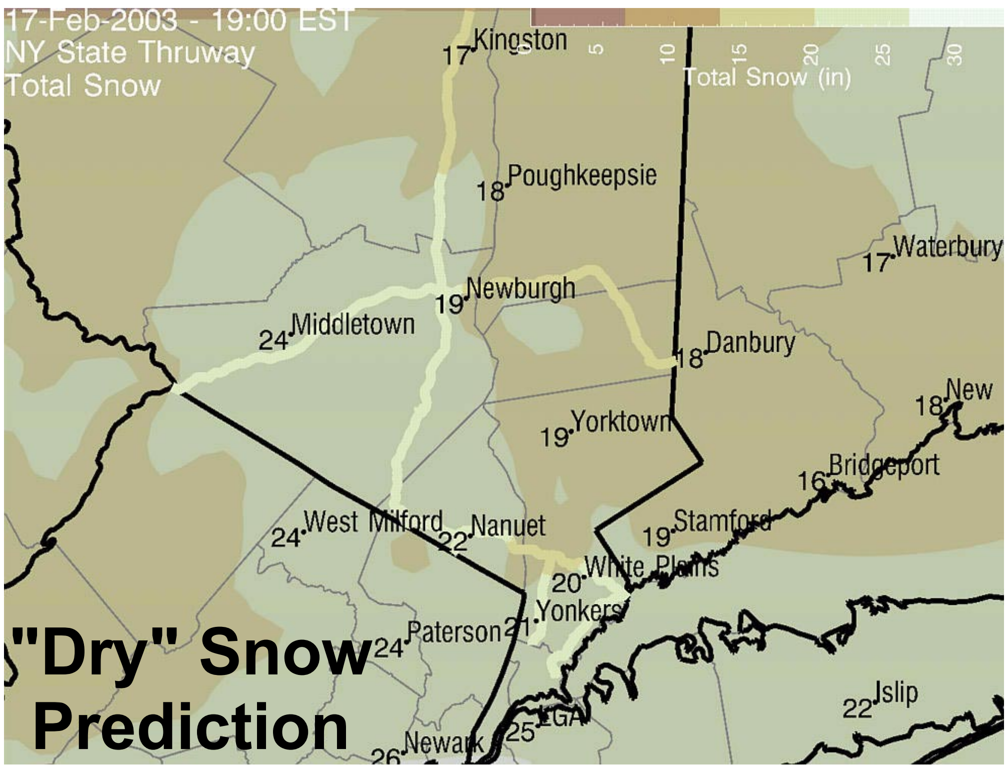

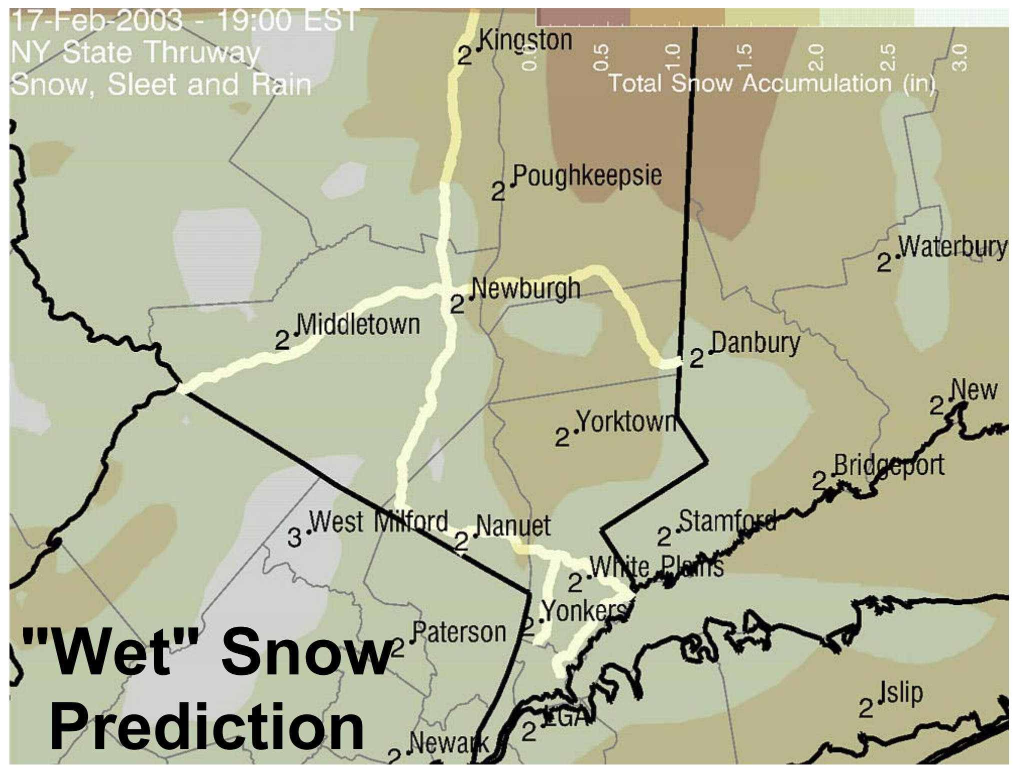

2/17/2003 Snowfall and Forecast*

Measured Snowfall (Inches)

Thornwood	26.0	Mamaroneck	18.0	Croton	14.5
White Plains	17.0	Yorktown	21.0	Yonkers	19.0

"Dry" Snow Prediction

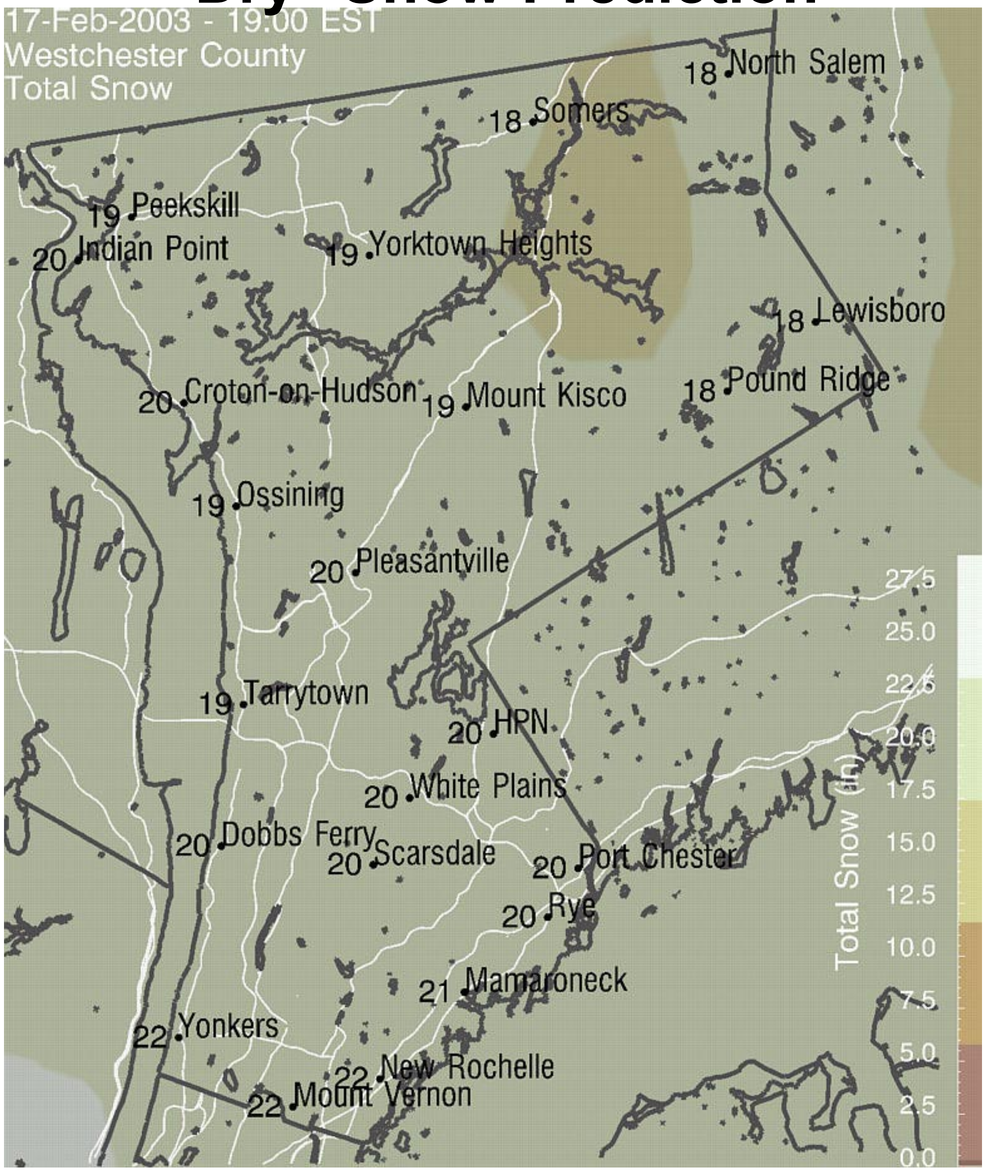

Liquid Equivalent Prediction

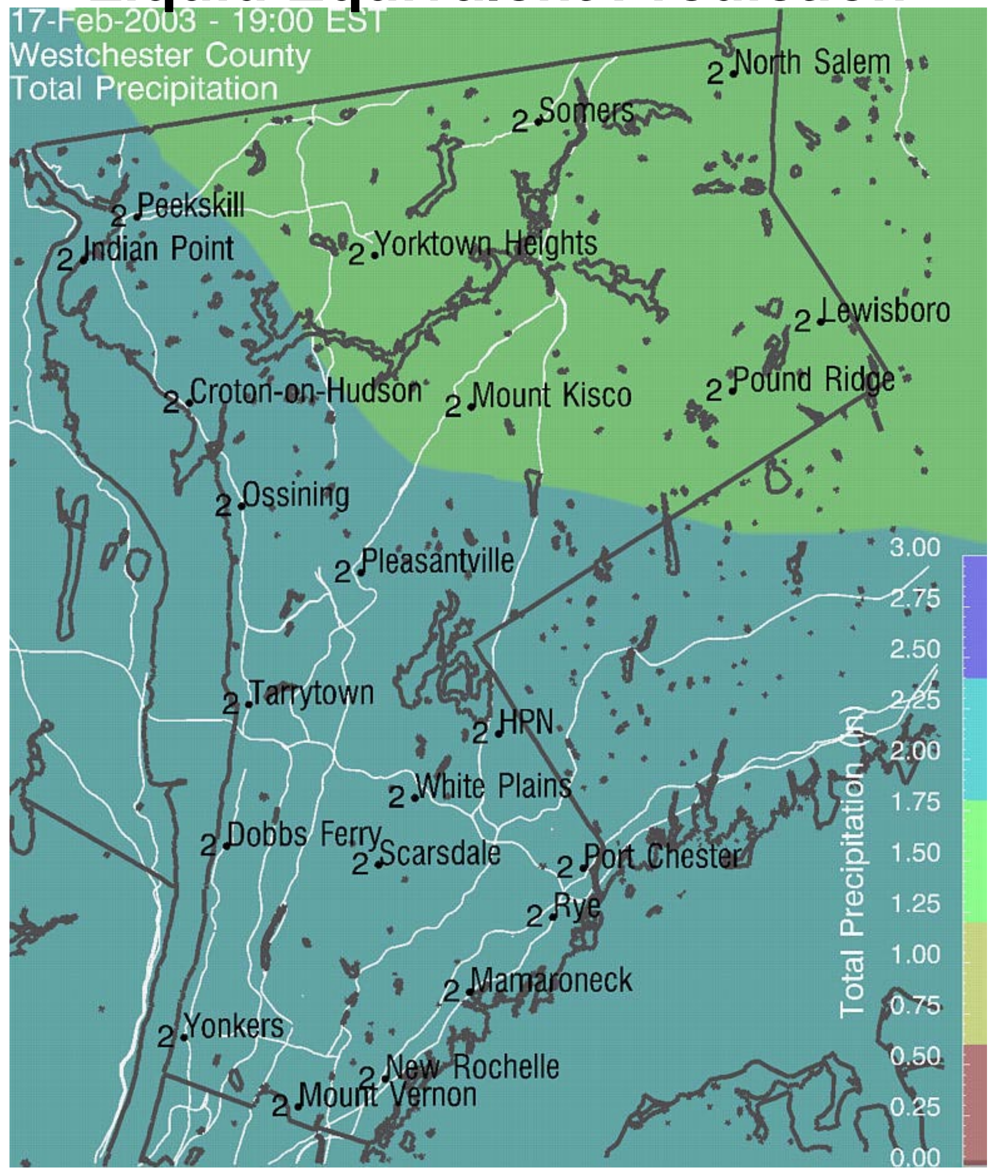

*Forecast initiated with data from 1900 EST on 2/16 with results available at about midnight on 2/17.

Example Web Products

Current Forecast

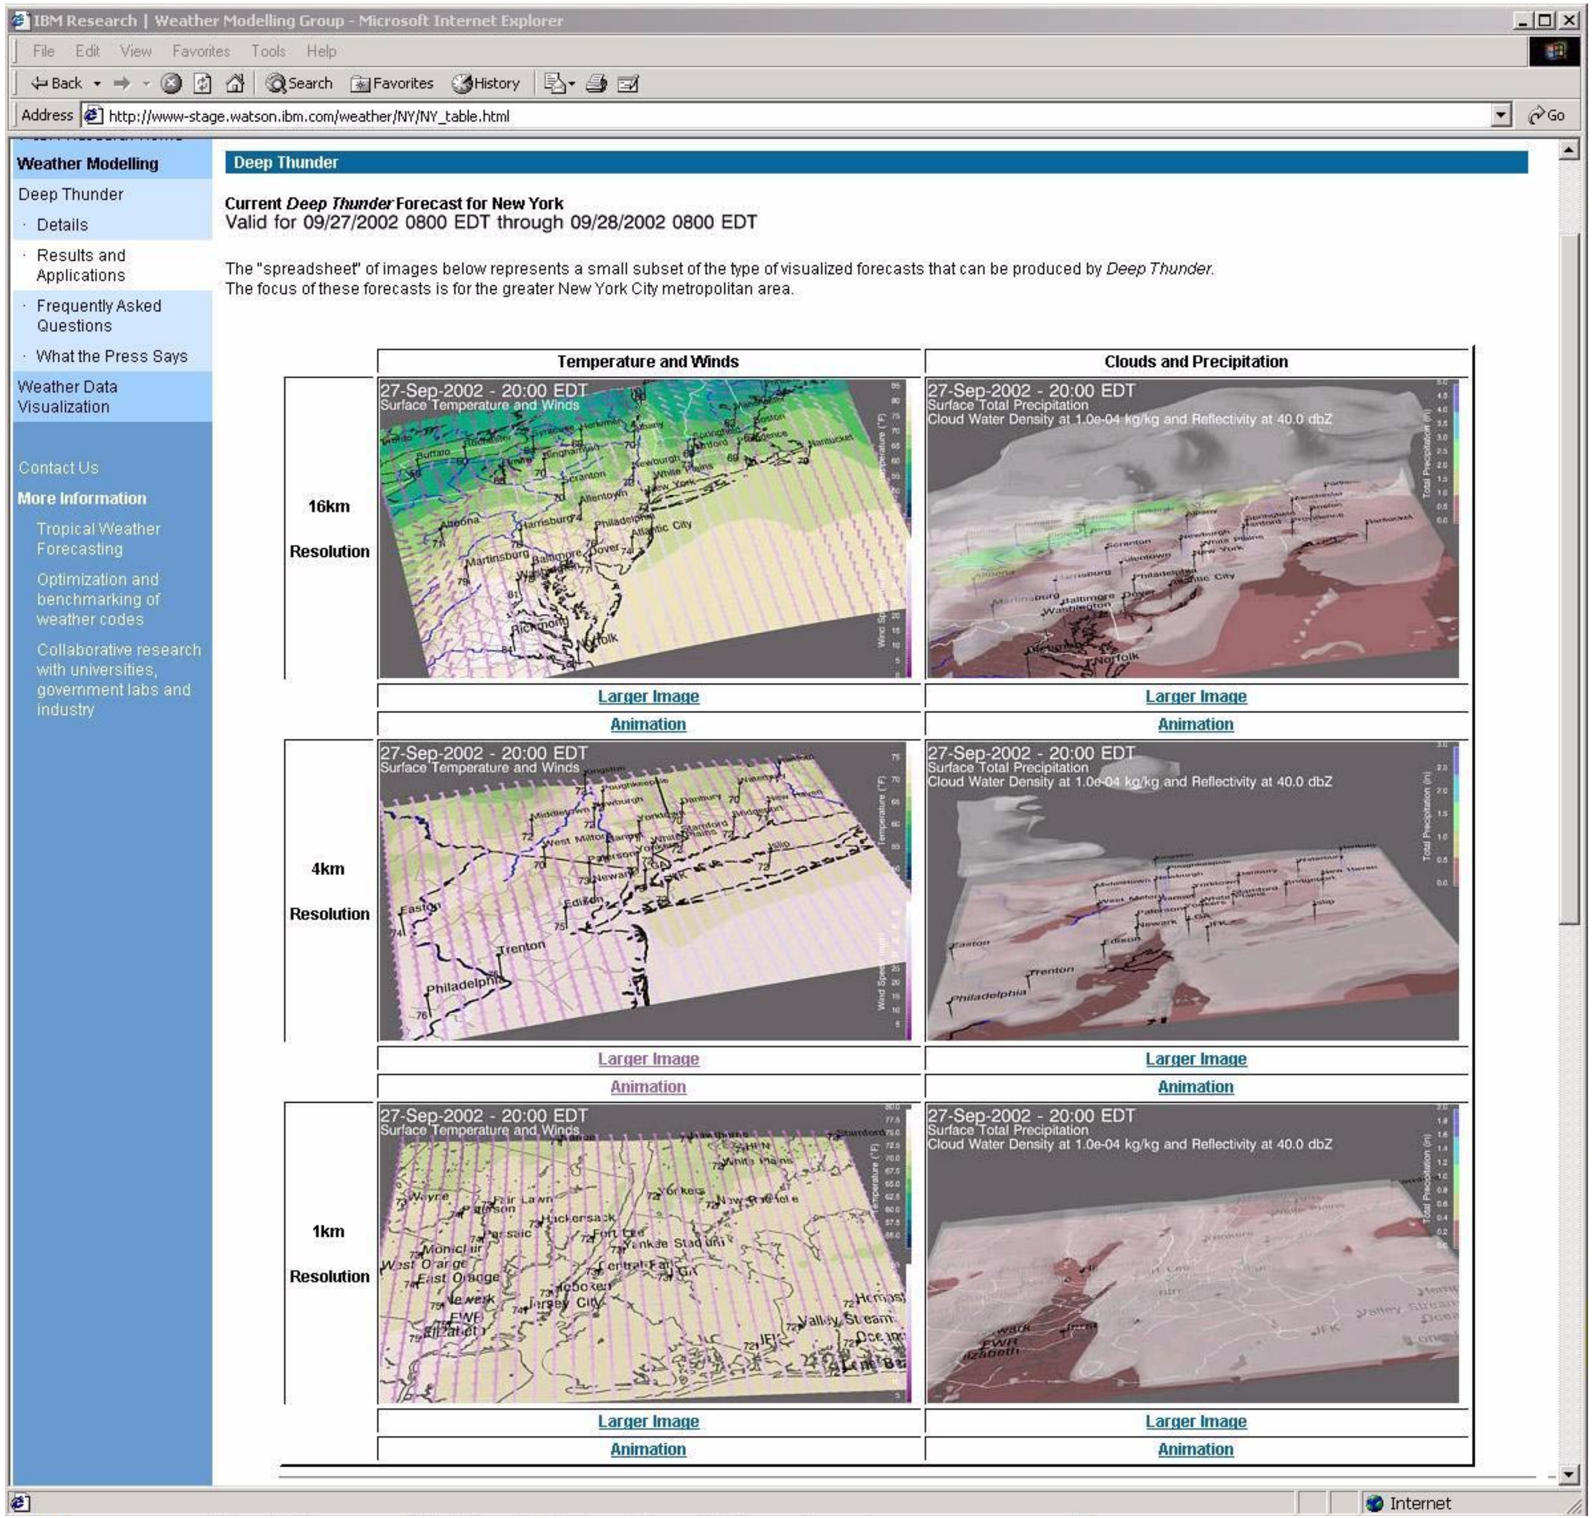

New York State Thruway Forecasted Road Conditions

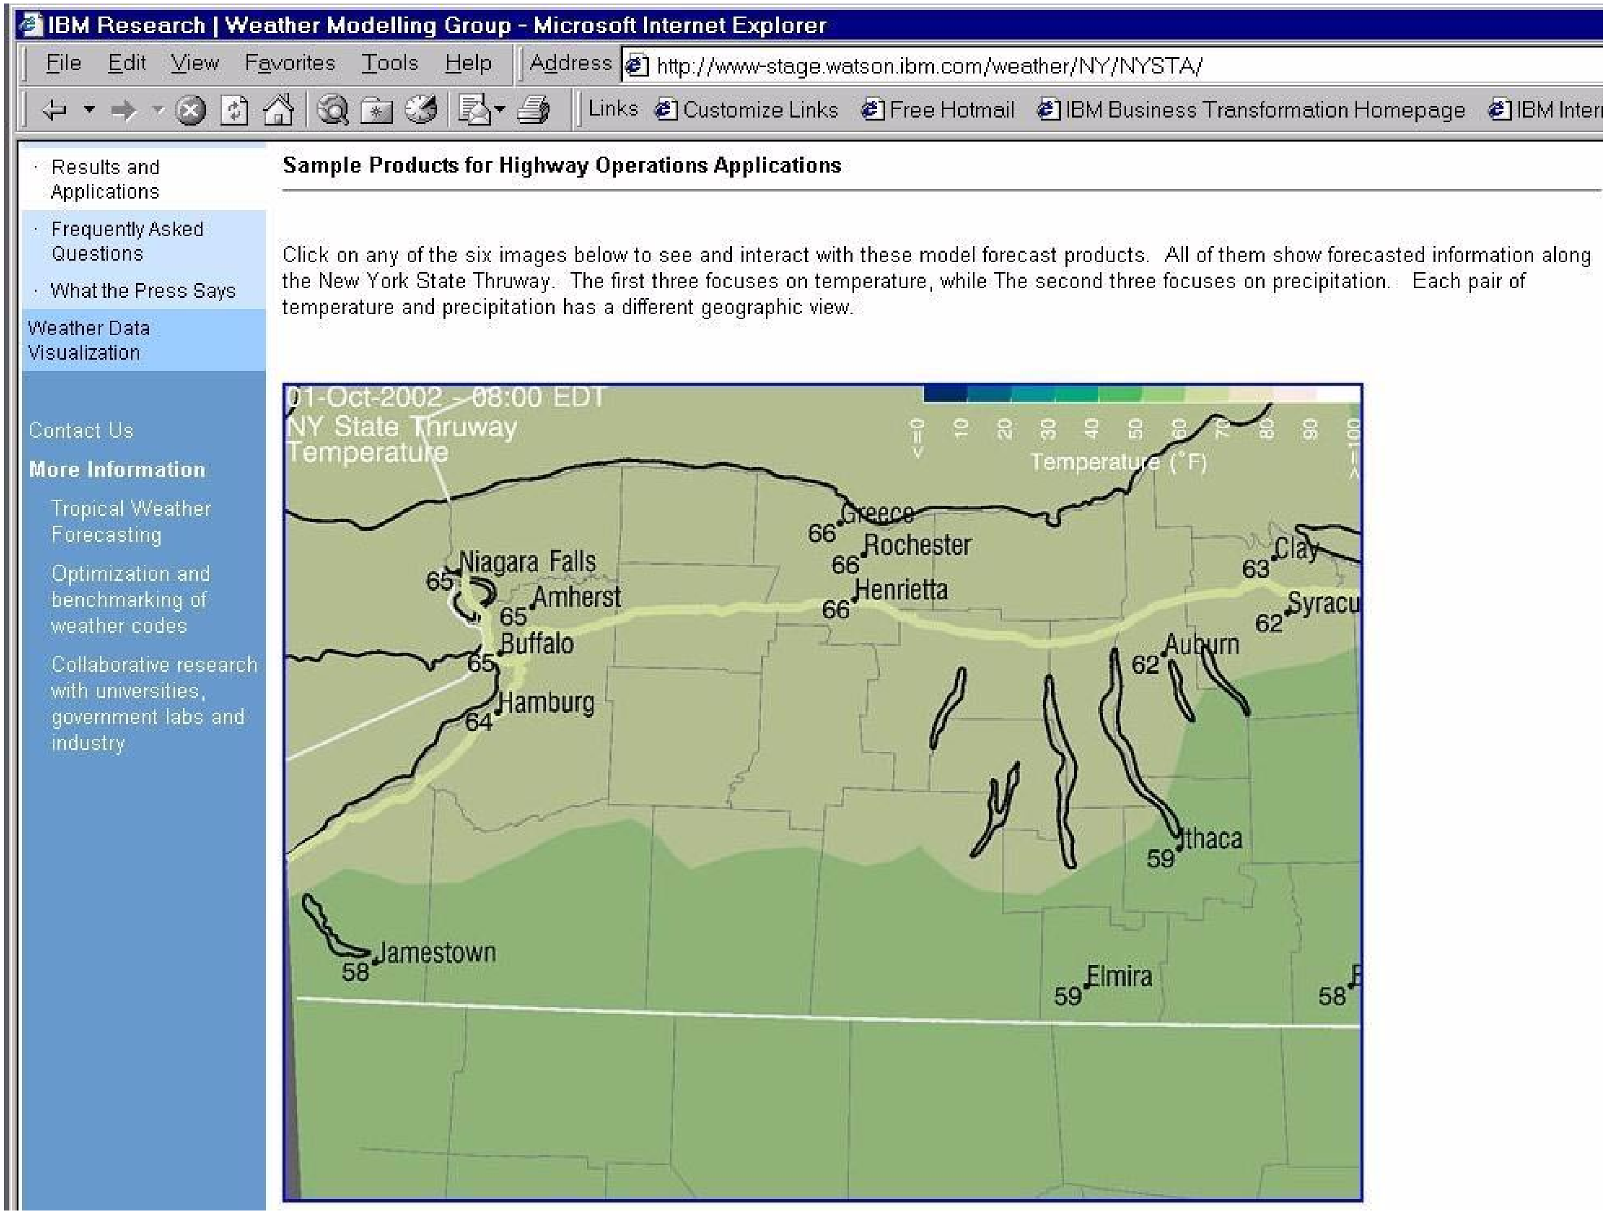

- Choice of weather variable and geographic region
- Color-contoured animated maps every hour

Westchester County Forecasted Conditions

- Choice of weather variable and geographic region
- Color-contoured animated maps every hour
- Location-specific surface and upper air data

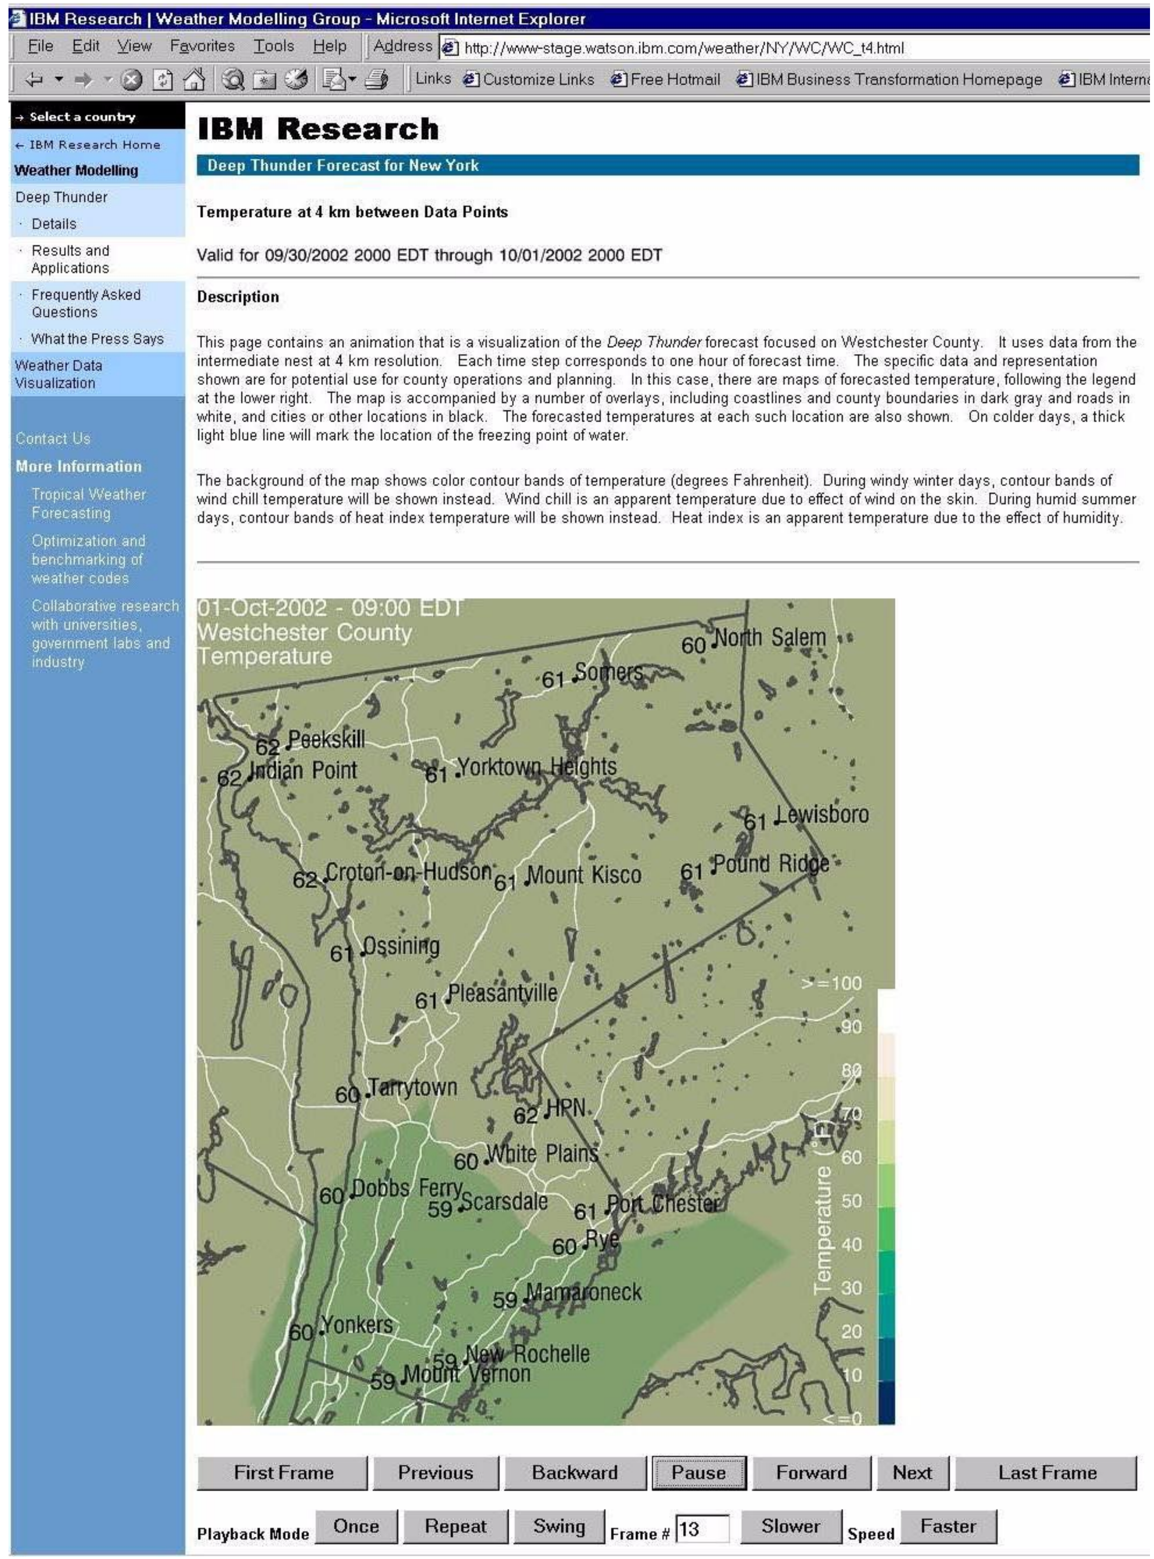

Aviation Flight Planning

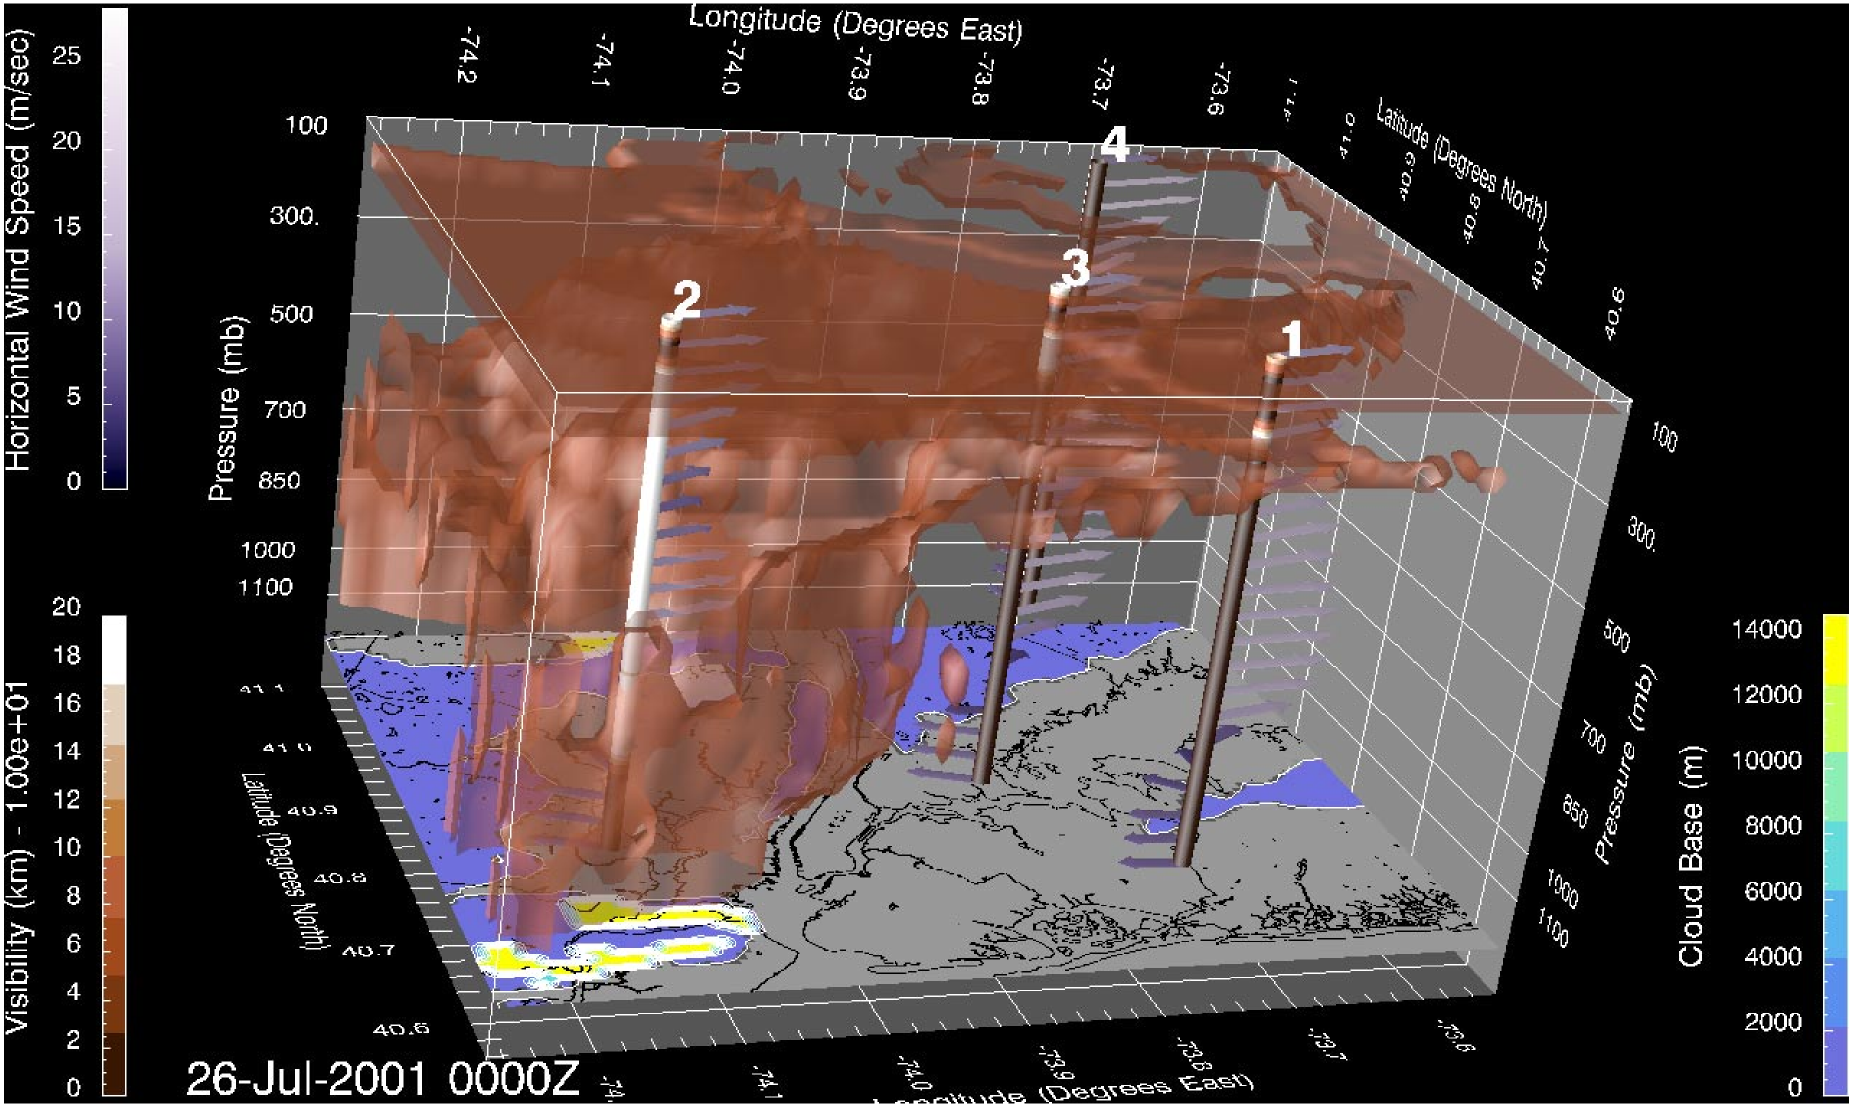

- **Visibility, clouds and winds**
  - Derived visibility: isosurface at 10km & profiles at airport locations
  - Contours (2km) of height of cloud base
  - Virtual wind profilers (speed and direction) at airport locations
- **Adaptation of analysis techniques for operational application**

Atmospheric Stability

- Identify areas for potential of severe weather
  - Color = K Index, convective potential based on vertical lapse rate along with amount and vertical extent of low-level moisture
  - Height = Lifted Index, potential instability from the surface to 500 mb
  - Contours = vertical wind speed
  - Green to yellow contours in blue "valleys" would imply regions of significant potential for severe convective activity.
  - Available in 1 hour intervals

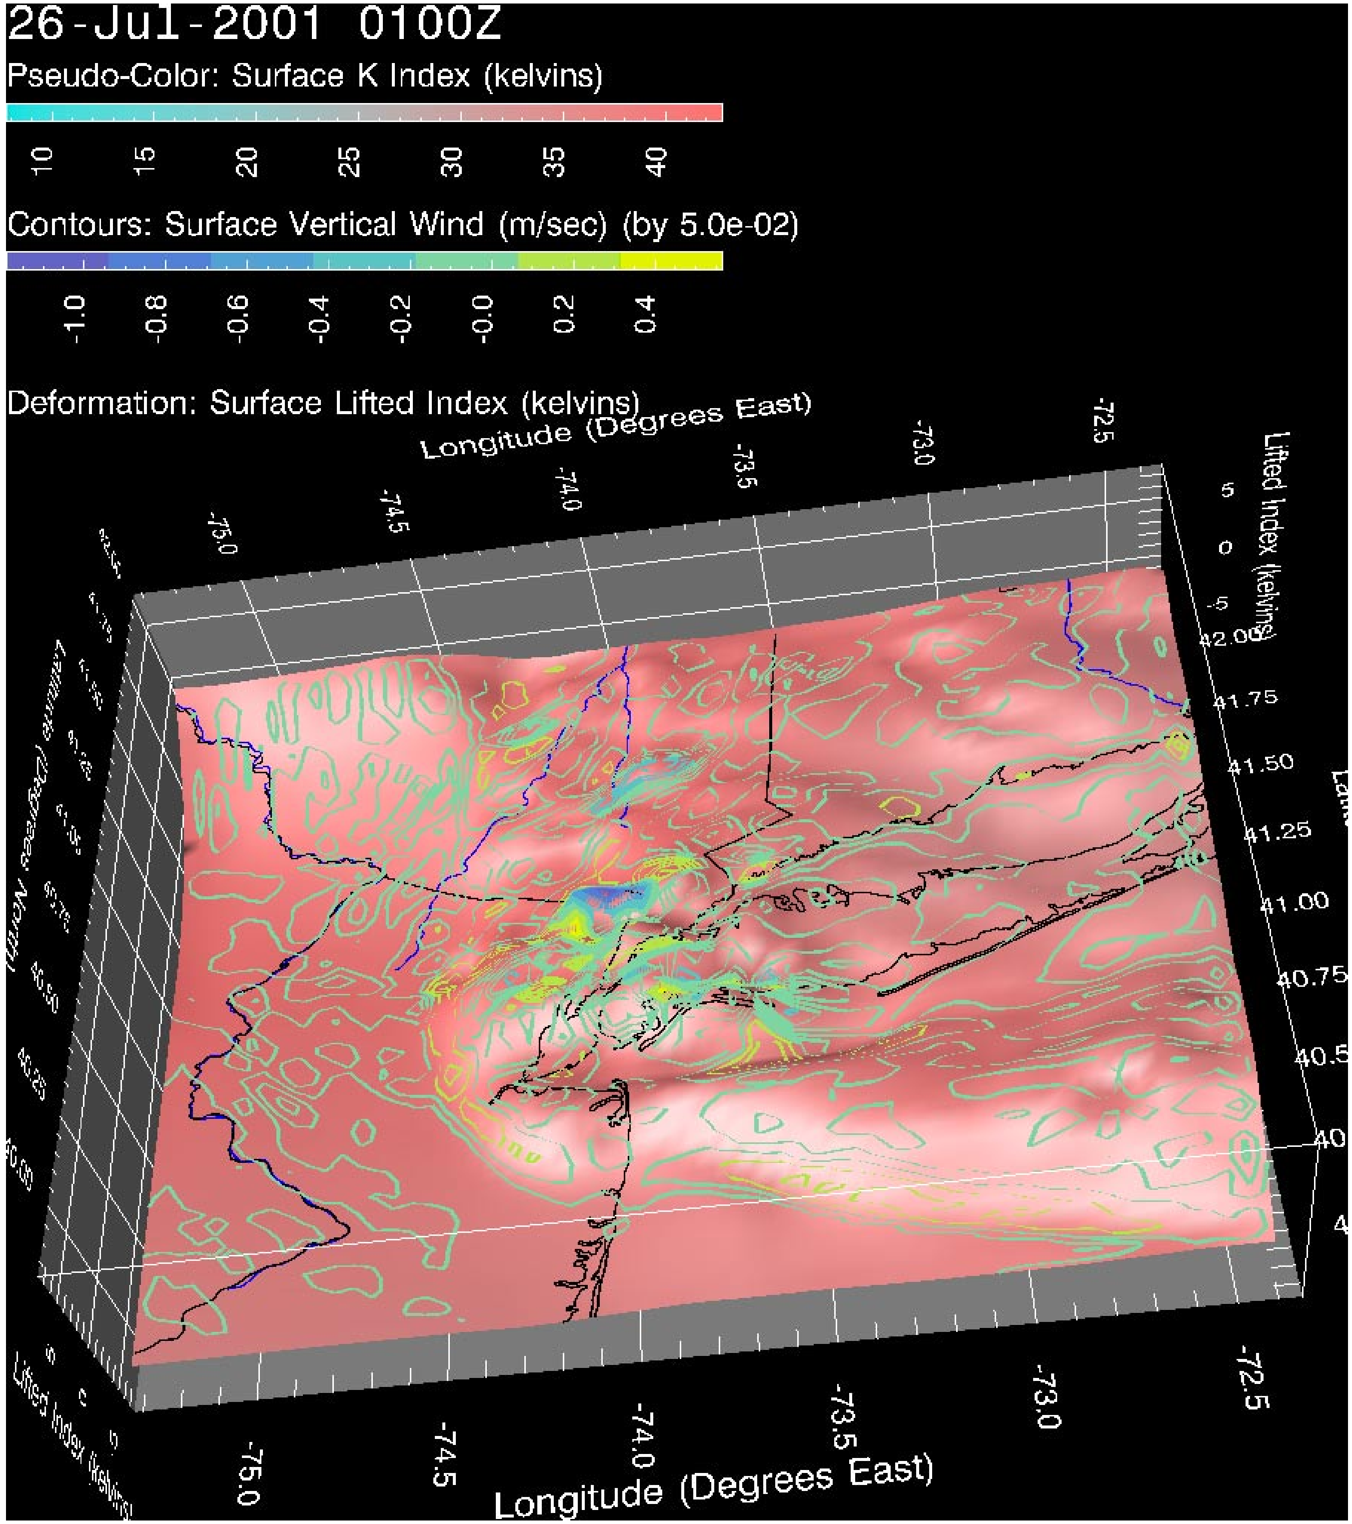

Example Snow Storm Forecast

07 February 2003

Site-Specific Forecast

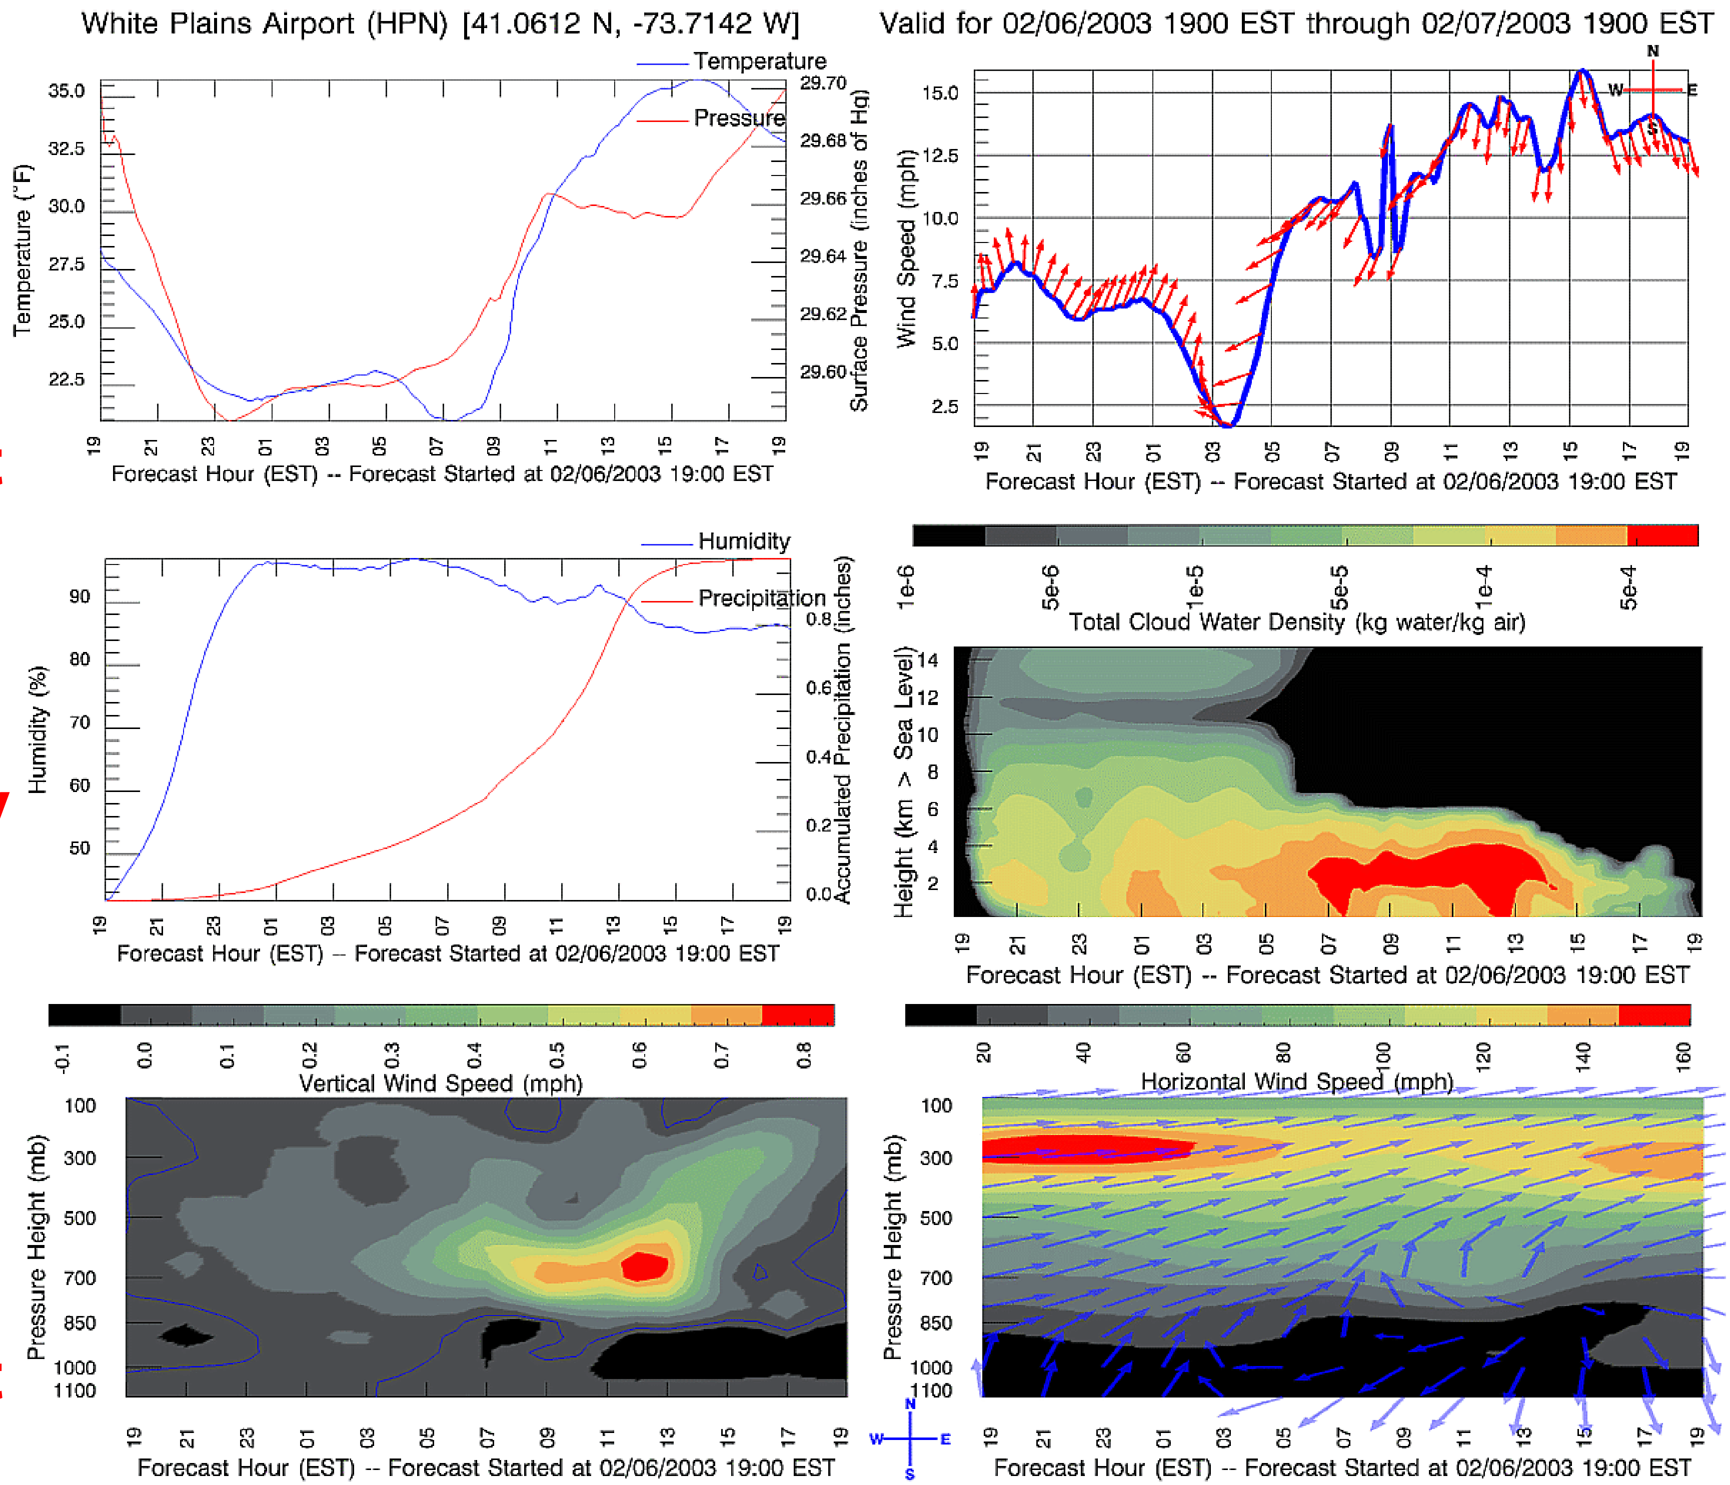

Deep Thunder -- Summary

- **An illustration of the viability of an integrated, cloud-scale, NWP-based forecasting system**
- **Observations by users**
  - Usable forecast produced are available automatically, in a timely, regular fashion
  - Favorable view of the ability to provide more precise forecasts of severe weather compared to conventional sources
  - Focused visualizations have been critical to effective utilization
  - But improved throughput and forecast quality is still needed
- **Continued research and development**
  - New York City prototype as platform for development and collaboration as well as evaluation of practical business applications
- **Future work**
  - Better computational performance and throughput coupled with additional model runs
  - Enhanced forecast quality and refined application-oriented product delivery
  - Deployment of small mesonet to aid in forecast verification and model tuning
  - Targeted verification (by area and application, e.g., travel delays, resource scheduling)
  - Improved initial and boundary conditions incorporating data assimilation

Thank You
Any Questions ?

Case Study Forecasts and the Visualization Tools Will be Illustrated During the Demonstration Session

Backup

Slides

Deep Thunder Modelling Component

- **Originated with the Regional Atmospheric Modeling System (RAMS), but extensively modified**
  - MPI-based parallelization for IBM Power/AIX SMP clusters, including nesting
  - Instrumented for visualization
  - Upgraded microphysics
  - 3-D, staggered in x-y, terrain following in z, moving grid
  - Arbitrary domain ($10\text{-}10^5$ m), nested, two-way interactive grid
- **Governing equations and numerical methods**
  - Unsteady primitive equations of motion, for all scales
  - Terms in equations are added/removed depending on scale
  - User selected physical and numerical schemes (e.g., split-explicit finite difference), both hydrostatic and non-hydrostatic
- **Operational configuration and performance**
  - Parallelization running 28-way (375 MHz Power3): 1.8 - 2.2 hours with current nesting
  - Initial and boundary conditions generated via isentropic objective analysis (θ as vertical coordinate) of relevant weather data

Nearest Neighbor Tool (NNT) Adapted from NOAA/FSL

Enables Model Parallelization on IBM SP

- Maps grids into logical grids of processors.
- Decomposes global arrays and assigns locally.
- Allows spaces for extra columns and rows.
- Assigns "I/O Caching" Processors (mpmd).
- I/O Cacher assembles data in memory.
- I/O is done asynchronously with "flush" commands.
- I/O and message passing are mapped onto MPI.

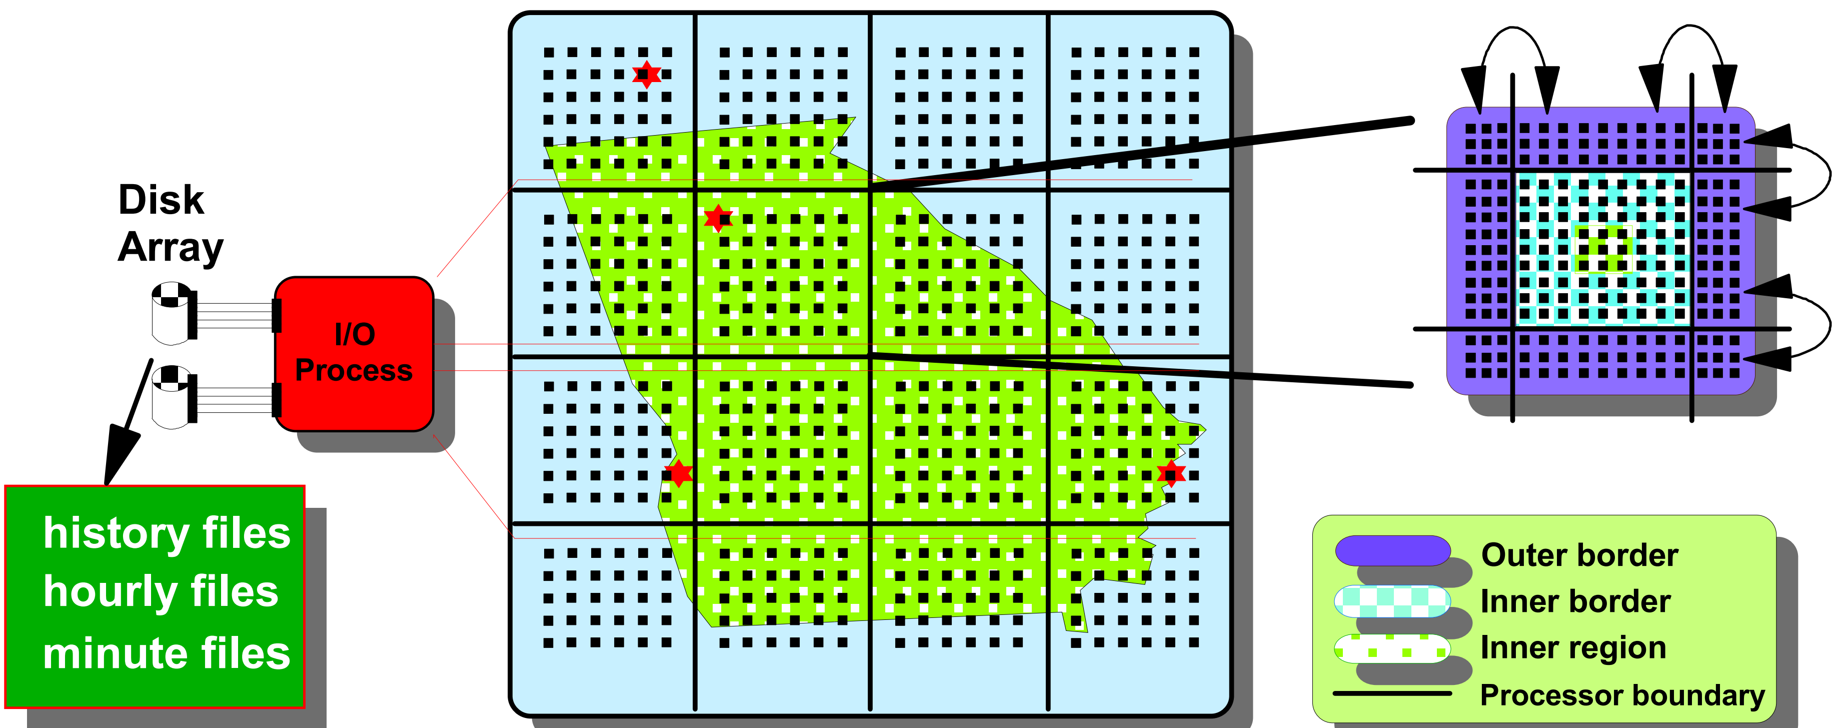

Forecast Processing Procedural and Data Flow

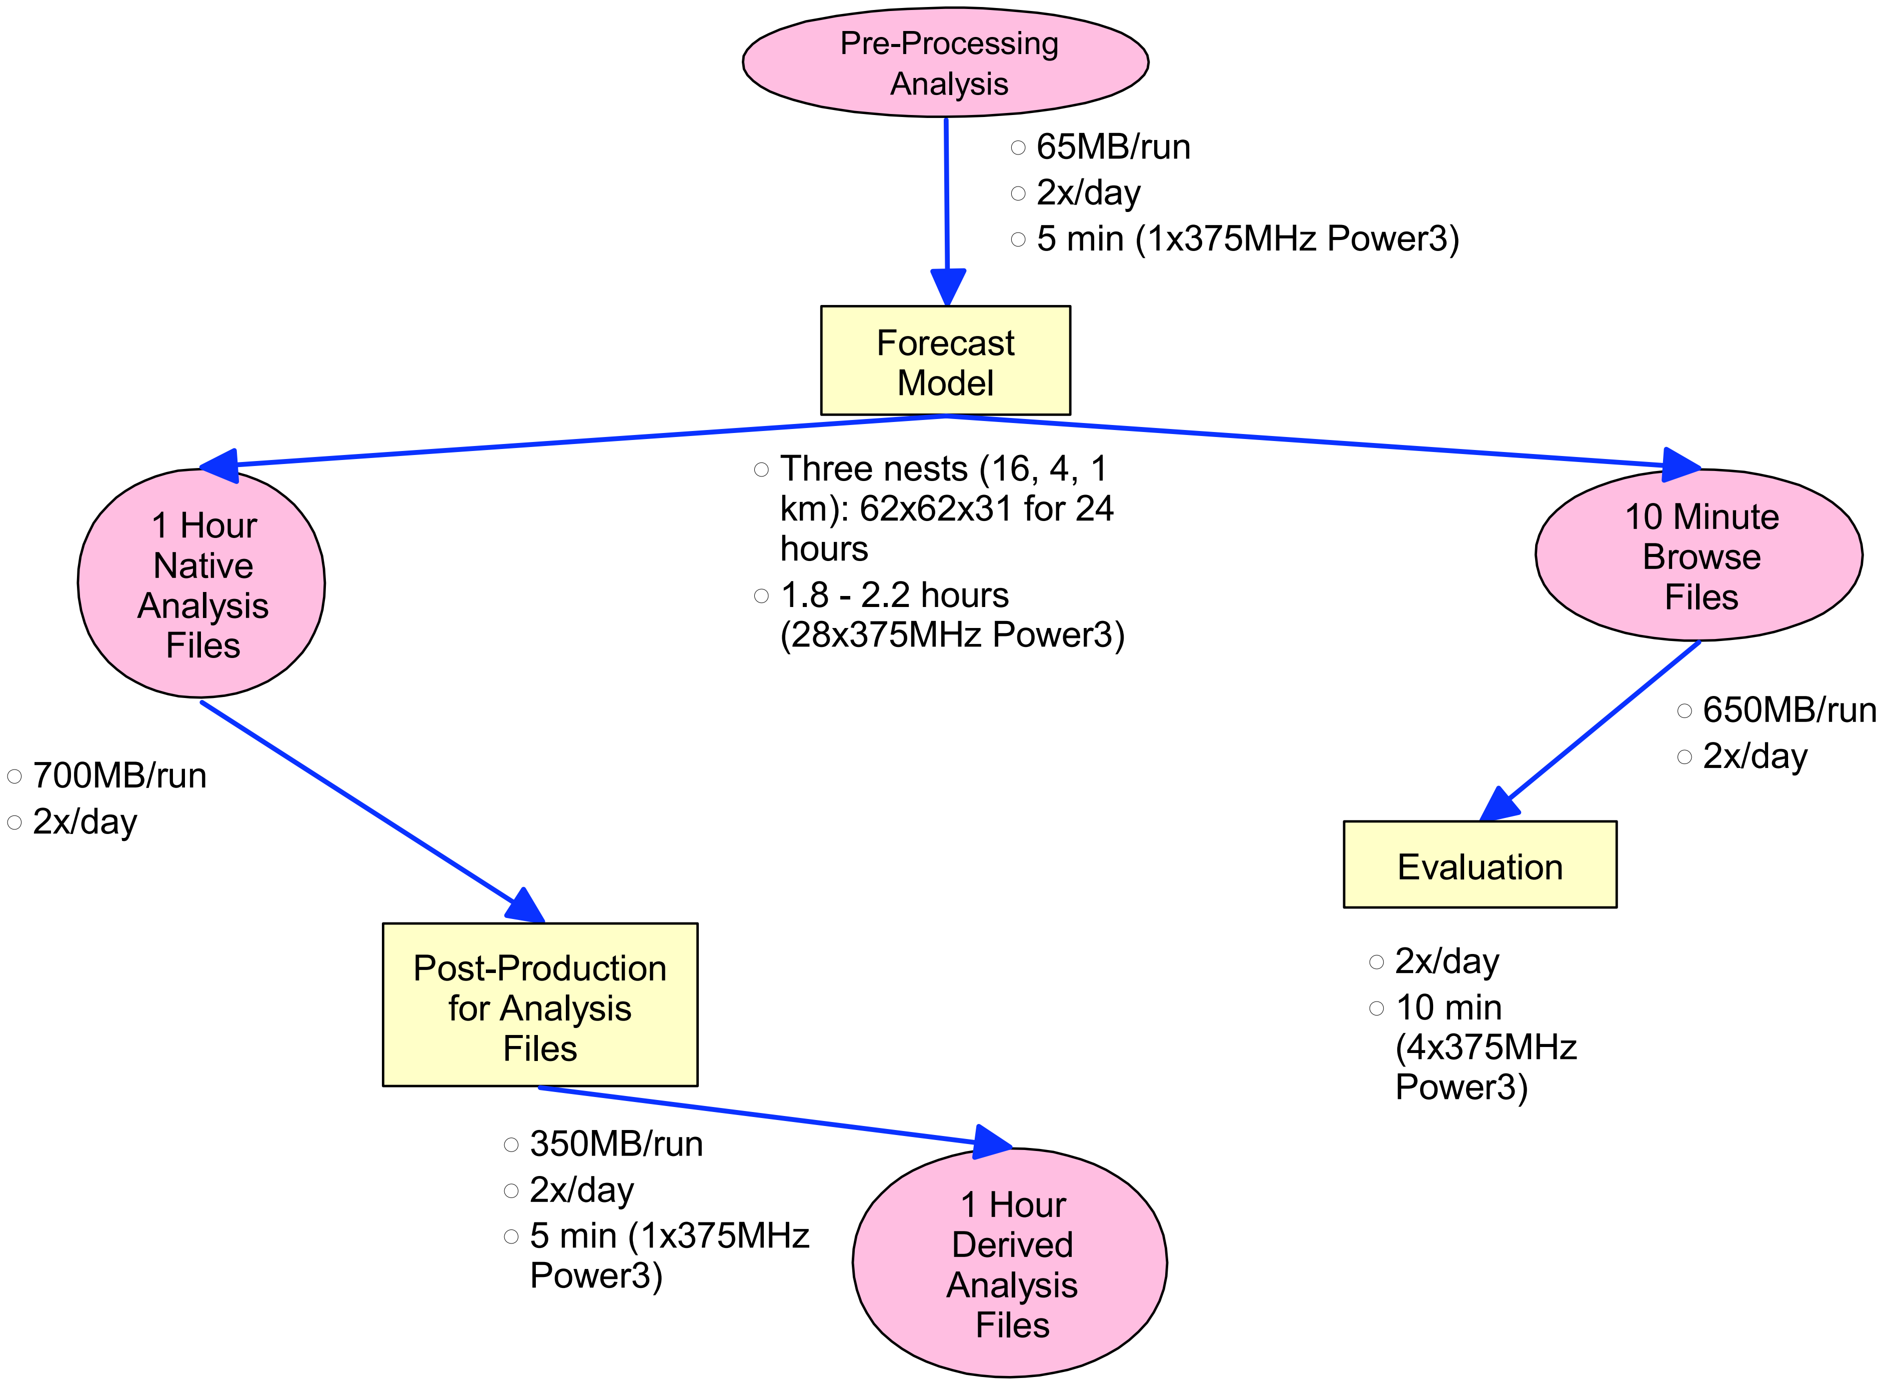

Visualization Implementation

- **Core implemented via Data Explorer -- an open source visualization toolkit (`www.opendx.org`)**
  - Custom tools for new visualization elements and derived meteorological variables
  - Custom tools/packaging for new output products
  - Shared tools and user interface components
  - Simple motif widgets for indirect interaction
  - Direct interaction in appropriate cartographic coordinates
  - *NO* transformation or compression of data or mesh(es)
  - Rule-based colormap tool used during design phase
  - Custom export/rendering www products
- **Integrated with mesoscale forecasting system (*Deep Thunder*)**
  - Custom I/O to balance communications, reduce latency
  - Custom filters for data import
- **Supplemented with utilities for image post-processing built upon ImageMagick (`www.imagemagick.org`)**

Lake Effect Snow *"Hindcast"*

- **Record "*Bald Eagle*" storm: 12/24/2001 - 01/01/2002**
  - Longest lasting lake effect snow on record
  - For example, 81.6" at Buffalo airport
  - Little wind and not particularly cold (daytime highs in 20s °F.)
  - Four parts of storm (two from Lake Erie, two from Lake Ontario)
- **Available archived input data for most of the event**
- **First test -- model forecast for second part of Lake Erie portion**
  - Last round of heavy snow, Friday, December 28, 2001
  - Bands moving south from southern Ontario and Niagara Peninsula
  - 0100 - 0700 EST: 15.2", 1100 - 1300 EST: 7"
- **Operational scenario**
  - Pre-processing and model execution begins between 2130-2200 EST (after input data received from the National Weather Service)
  - Model execution and initial post-processing completed between 2300-2330 EST
  - Initial products and then analysis available between 2300 and 0000 EST

December 27 - 28, 2001 Lake Effect Snow Hindcast

- Air temperature below freezing
- Snow band already moved north

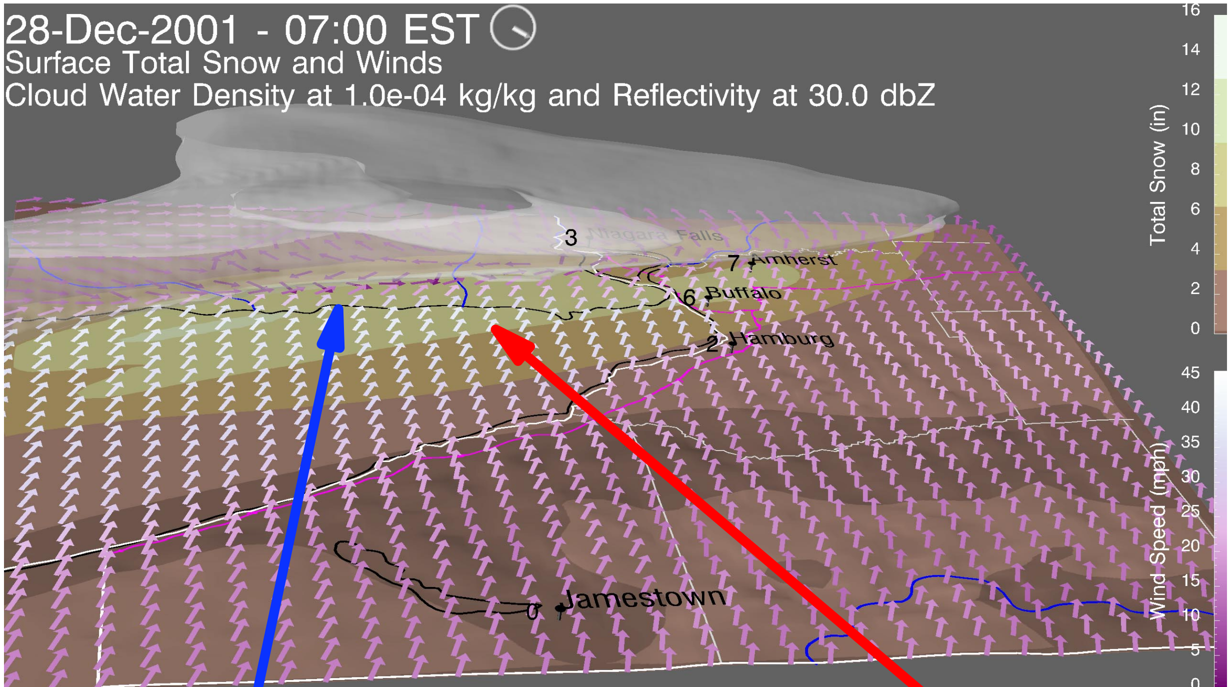

- 18, 6 and 2 km grids (68x68)
- 24-hour run initialized at 0 UTC on 12/28/01
- Small negative bias on total snow amounts

Winds shift direction

Lake Erie

- Air temperature rises above freezing over Lake Erie
- Snow band moves south from Ontario and Niagara

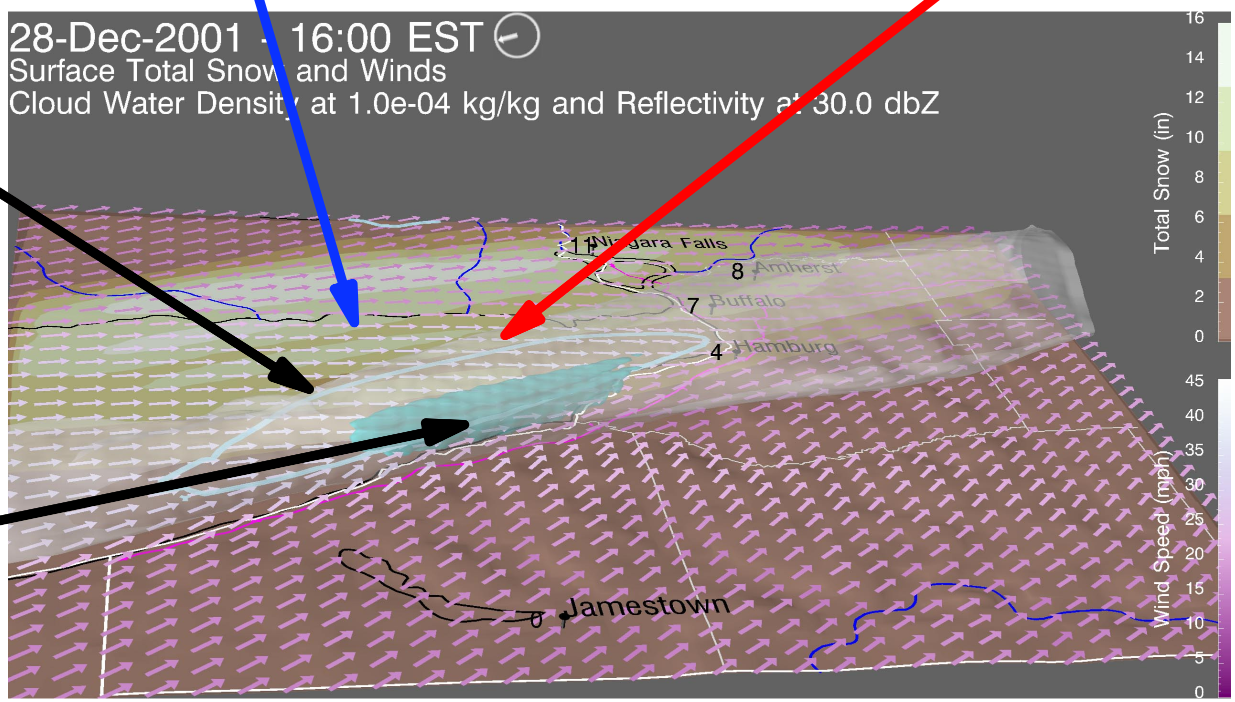

- Positive bias on timing of about one hour
- Lake-effect snow band resolved
- Reversal of wind direction forecast

Different Visualizations of a Thunderstorm Forecast

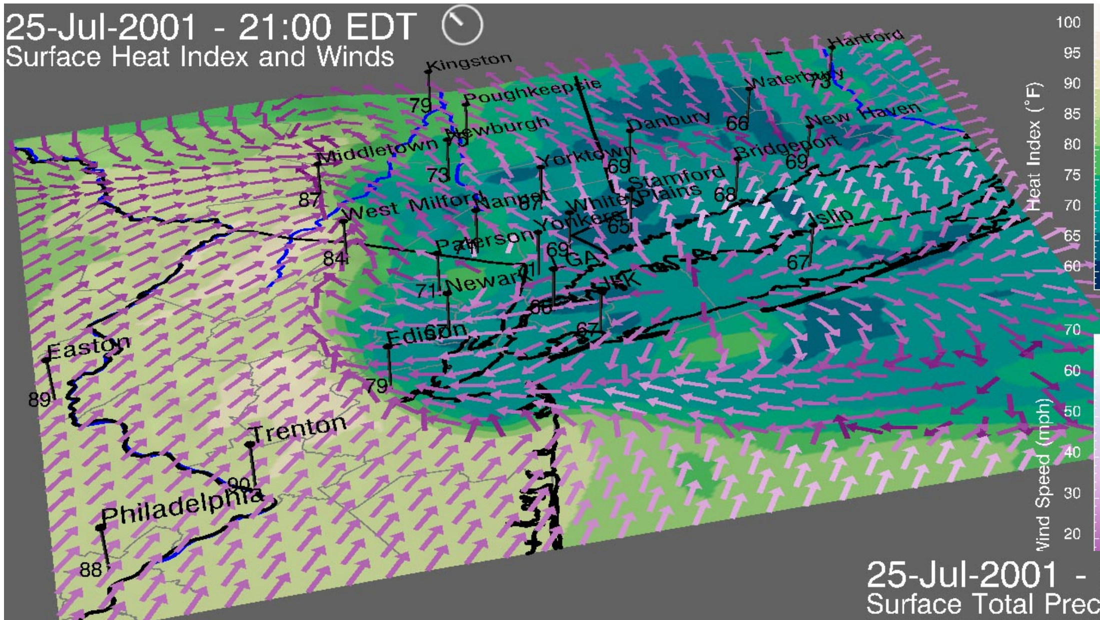

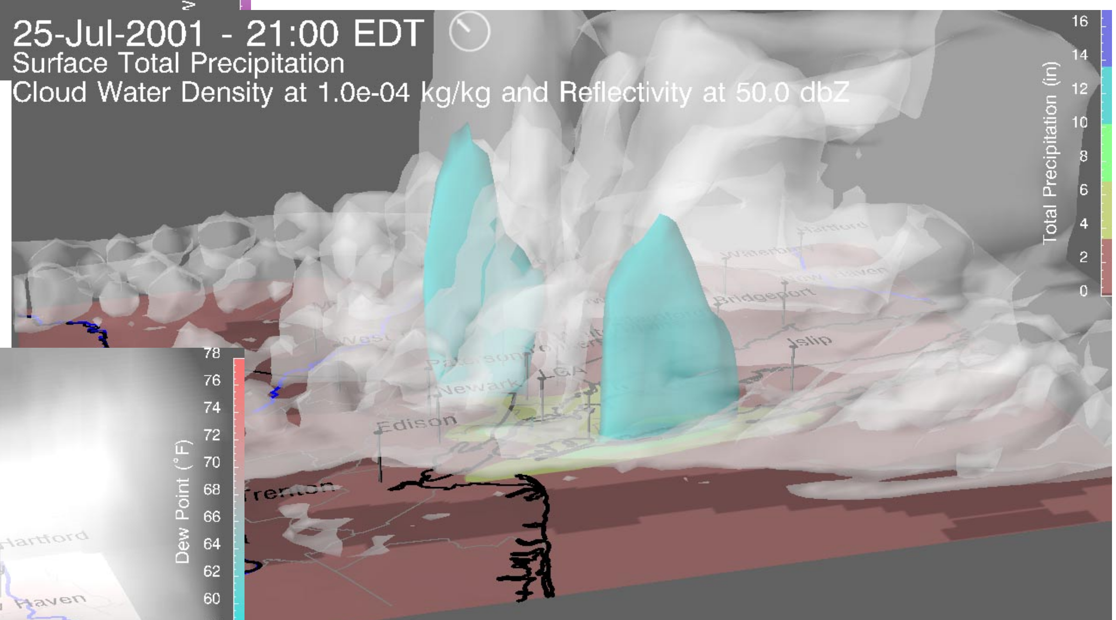

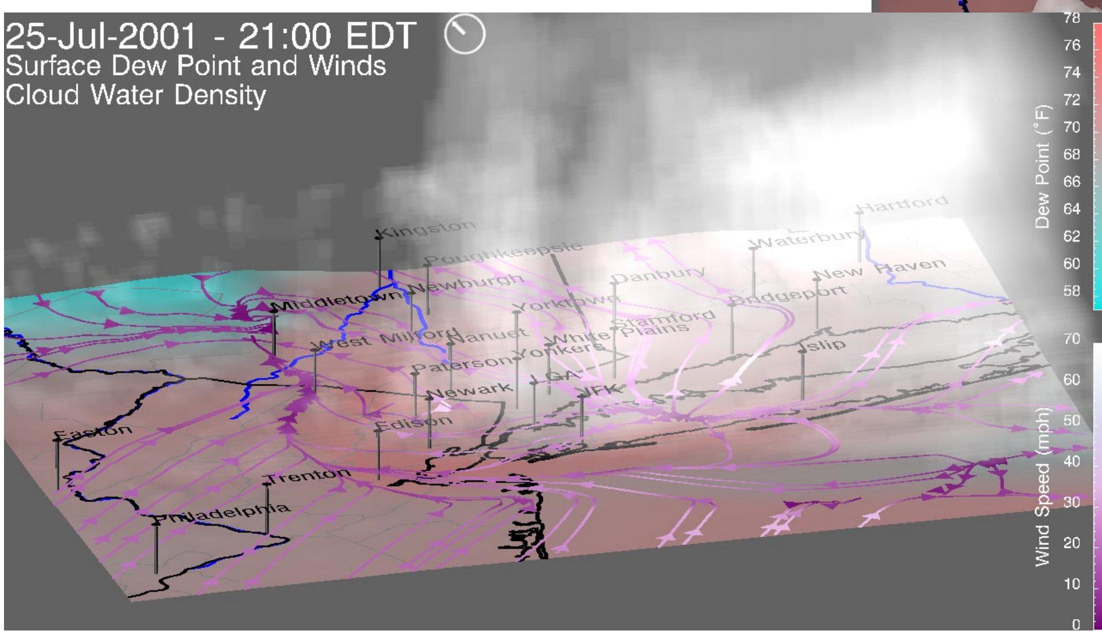

- Surface wind, heat index, dew point temperature, total precipitation
- Three-dimensional cloud water density and reflectivity
- 4 km resolution nest

Quality of Current Operational Forecasts

- **Very good skill at forecasting severe or unusual events**
  - Consistently better snow forecasts for the 2002-2003 winter than what was available, coupled with far greater detail and longer lead times
  - Predicted ~90% of severe weather recorded in 2002 that impacted NYC airport operations
  - Predicted N80% of thunderstorms in 2002 within 2 hour time window
- **Forecasts are provided with significantly greater information compared to others**
  - For example, timing, location and intensity vs. "chance of thunderstorms in the area"
  - Far longer lead time compared to weather warnings based upon current observations
- **Compared to a limited set of surface observations, accuracy is comparable to other results**
  - For temperature, humidity, winds, etc.
  - Yet still provides additional benefits via precision and packaging