



## Predictive Analytics and Dynamic Thresholds

### Current Capabilities and Futures

*David Curbishley, Service Management Solution Architect, Tivoli Software, IBM*



# Disclaimer

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.



# Agenda

- Why Predictive Analytics?
- Current ITM Challenges
- How is ITM 6 Roadmap Impacted
- What Opportunities Result?



# Predictive Analytics is a Response...

## ...to Customer Requests

- I'd like 30 minutes warning to know when my user experience is going to deteriorate.
- What are realistic baselines for my environment?
- I'd like to adjust for repeating traffic situations so I don't receive false alerts.
- I'd like my operators to be able to more quickly diagnose certain events and patterns to implement fixes.

## ...to the Outage Avoidance opportunity

- Moving IT from reactive to proactive/predictive.
- CIOs want their business' to benefit from identifiable outage avoidances
- Helps show what IT is saving and delivering versus just being a cost center.



# Extending Tivoli Suite with Predictive Analytics

**Built in predictive analytics & value at all management layers -- not *an overlay product!***

- ✓ Predictive service impact, RCA, SLA tracking & event/performance mgmt.
- ✓ Broad predictive collection & experience across: mainframe, power, virtual, SOA...
- ✓ Common warehouse & Visual, Navigation, SSO, Process automation
- ✓ Broad integrations across IBM & 3<sup>rd</sup> party collection

- ✓ Extended analytics for metric forecasting
- ✓ Additional domains including security and storage

## Service Availability & Performance Management



### Roadmap to Predictive Leadership

- ✓ Robust CMDB integration
- ✓ Run Book Automation
- ✓ Extended 3rd party agent library
- ✓ Performance Management Database

### Multi-layer Approach

- ✓ Assures predictive value at any phase
- ✓ Scales to any size environment
- ✓ Maximizes OOTB Intel, while minimizing collection of irrelevant data
- ✓ Leverages and improves ROI on existing investments

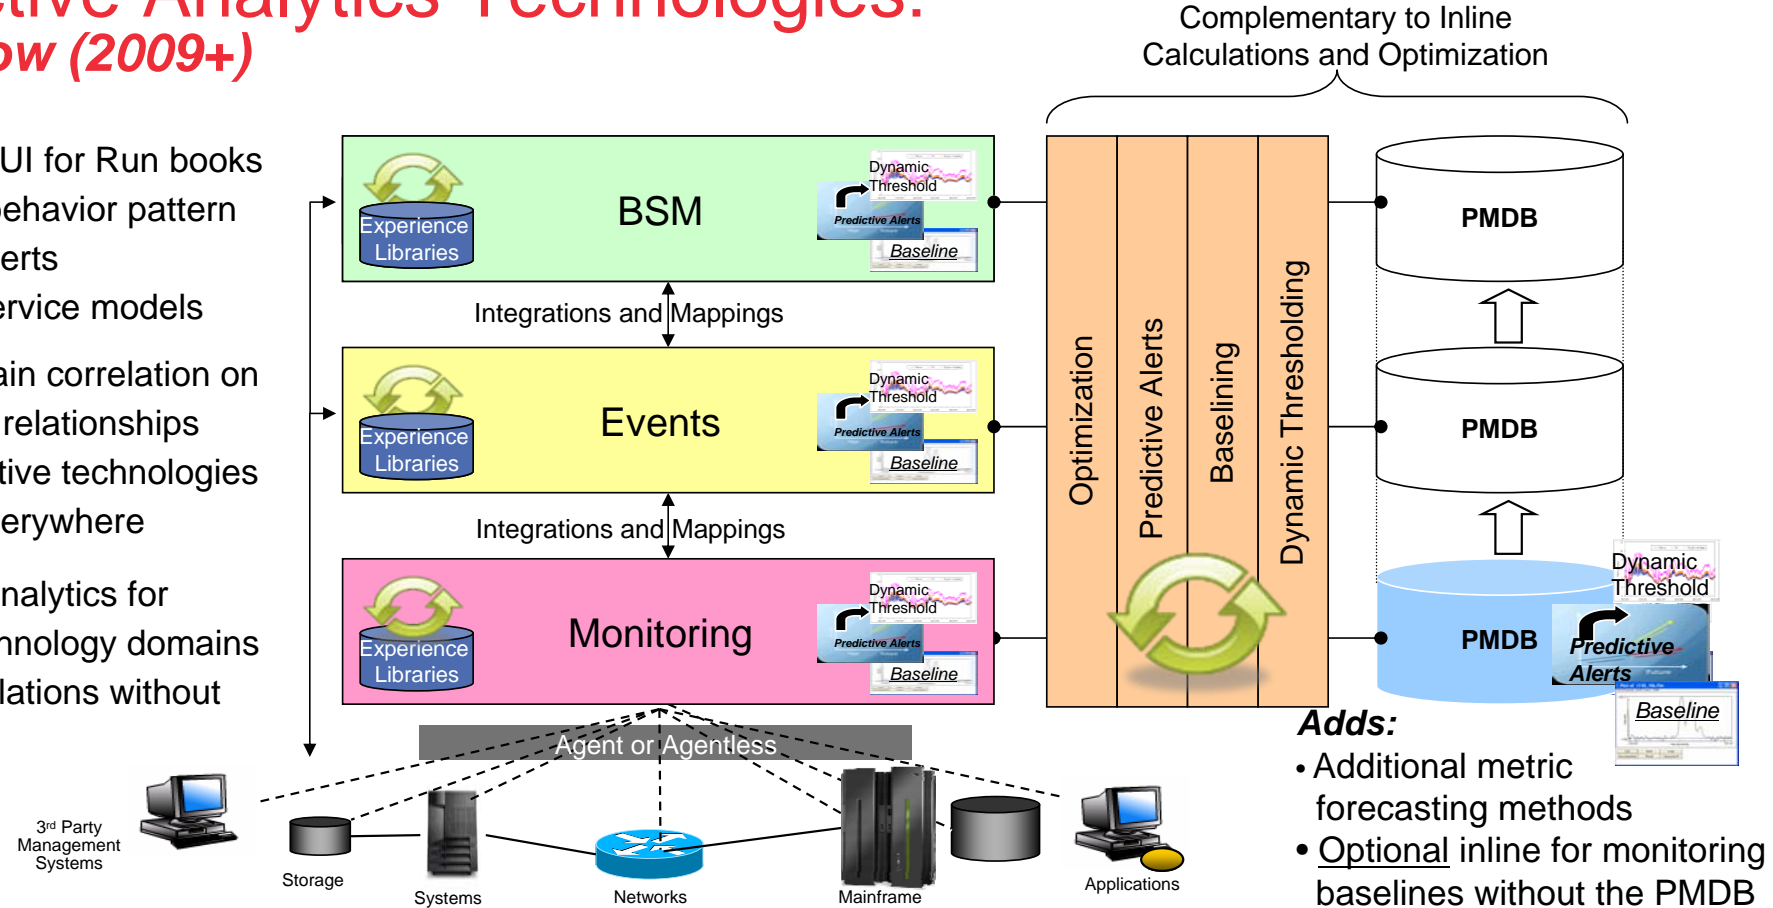


# Predictive Analytics Technologies: Tomorrow (2009+)

## Adds:

- Workflow GUI for Run books
- Abnormal behavior pattern detection/alerts
- Adaptive service models
- Cross domain correlation on discovered relationships
- Inline adaptive technologies included everywhere
- Additional analytics for specific technology domains
- Inline calculations without PMDB

Out-of-the-Box Integrations




## Adds:

- Additional metric forecasting methods
- Optional inline for monitoring baselines without the PMDB




# Predictive Analytics: A New Approach to Deliver Outage Avoidance

Evolve Your Approach to Management



IT Operations Predictive Events Histogram



	State	Events	Predictive Status
401K Monte Carlo	▲	▲	
Systems	▲	▲	▲
CPU	▲	▲	▲
Disk	●	●	●
Memory	●	●	●
Network	●	●	●
401K Transactions	●	●	●


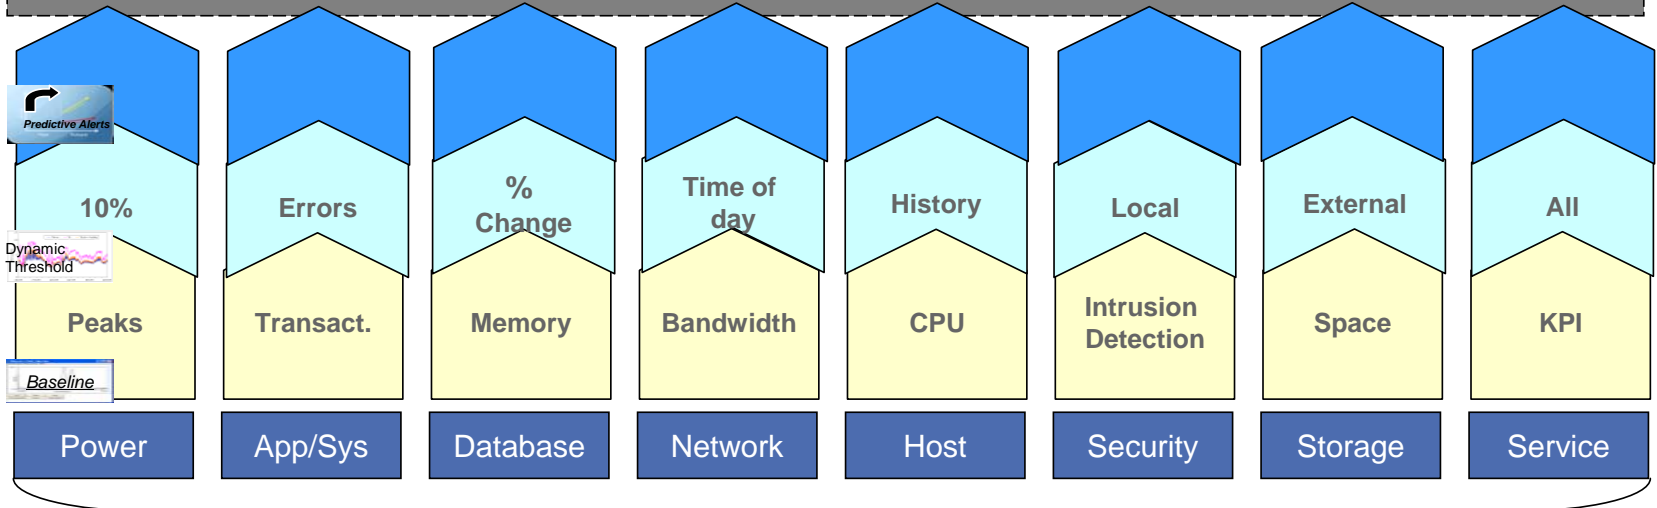
Ability to aggregate or focus by domain/resource, local need, schedule, history, or external.

Predictive Alerting

Dynamic Threshold

Adaptive Baselines

Resource



# Advanced Monitoring Analytics: IBM Tivoli Performance Analyser (ITPA)

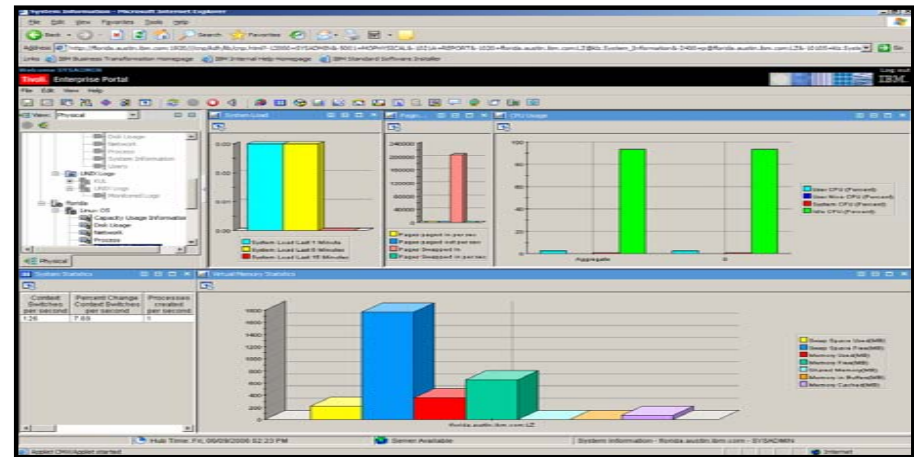
## What ITPA Does

- Provides hands off capacity monitoring
  - Automates performance analysis and reporting
- Enables prediction of application bottlenecks and creation of alerts for potential service threats.

## Supported Scenarios

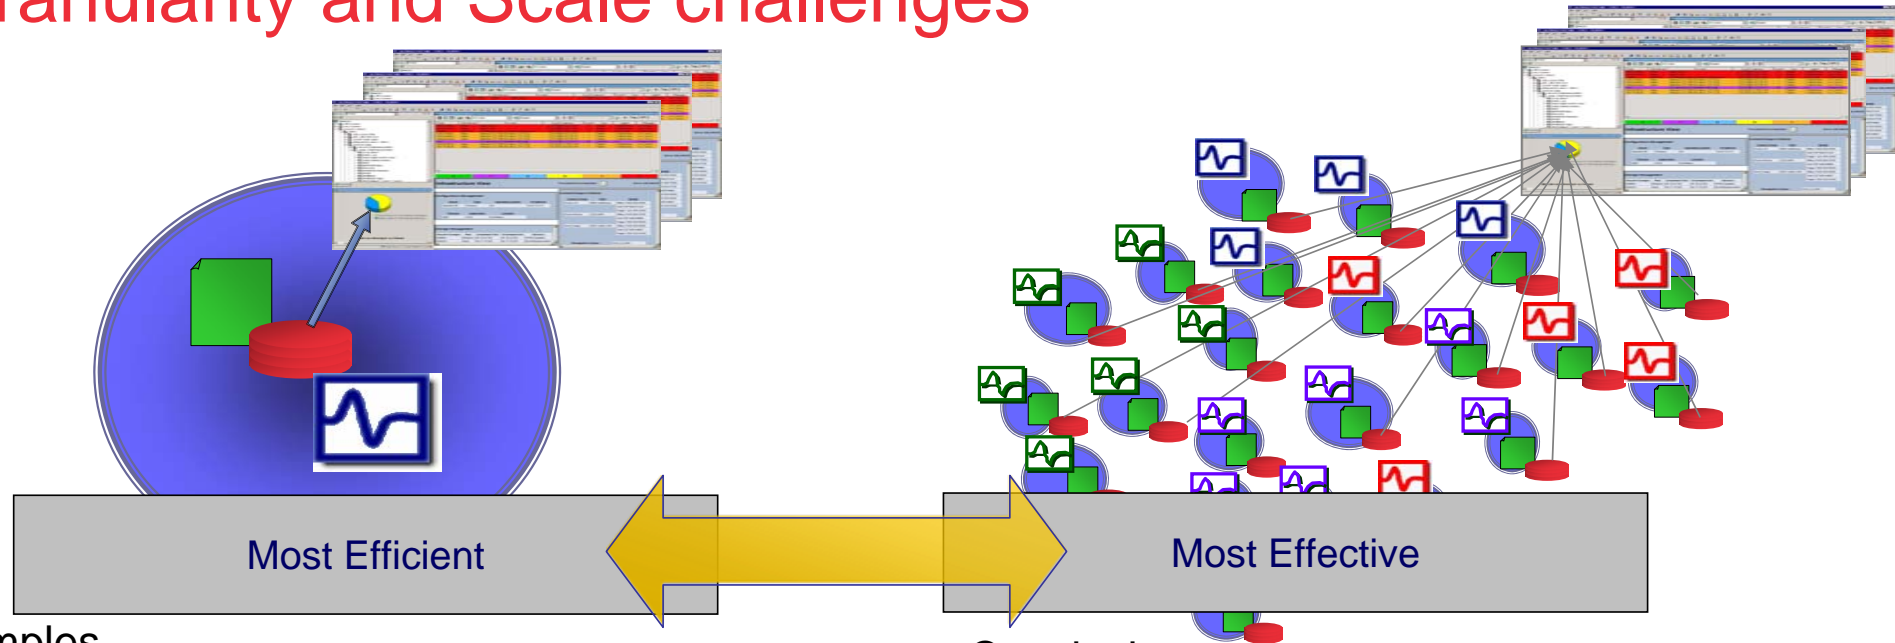
- “What will my resources look like tomorrow, next week, next month or next year?”
- “What IT Resources should I worry about next?”
- “Will I have enough capacity to get me through Monday?”

- A seamless extension to ITM 6
  - Uses existing ITM agents
- Derive new metrics using arithmetic expressions
- Predictive trending and forecast reports
- Out of the box reports for distributed systems
- Fully Extensible
  - Supports the Universal Agent and Agent Builder





# Granularity and Scale challenges



## Examples

- Generalist IT Manager vs. Domain Specialists
- One set of thresholds for all servers vs. thresholds per server
- Centralized event forwarding vs. distributed eventing
- One set of authorizations vs. authorizations per user
- One UI for the admins vs. one for each admin

## Conclusions

- Data, Control and Thresholds have similar range of requirements
- Granularity of performance data - collected and retained
- Visualization – Who and What?
- Security



# The Required Range of Configuration Options

- Not an ITM-centric perspective – Portfolio-wide view
- No “one size fits all” options
  - They vary:
    - Across industries
    - Across customer sizes
    - Within a single customer
  - Best option is a trade between service level requirement and affordability
- Options must be independently customizable
- Minimize disruption as requirements change



# 2008 ITM Releases

## ITM v6.2 FP1 – GA: Apr 2008

- Cache Pure Event Results
- Single Sign-On and ESS Support
- Firefox support
- Sun JRE 1.5 Browser Support
- ★ CLI
  - Consistent cross-agent versioning from CLI
  - Maintenance mode
  - Take Action
  - User & Group Management
  - Remove off-line agents from tree
  - Bulk export of situations preserves distribution lists
- Infrastructure / Agents
  - MS SQL 2005 Support
  - Windows Vista Agent Support
  - Use of localized gskit for install
  - i5/os agent enhancements – inactive job monitoring, support for new releases, detecting when user profiles are disabled, cross site mirroring, additional work spaces

★ Denotes granularity/scale impact

## ITM v6.2.1 – GA: Nov 2008

- ★ Dynamic Thresholding
- ★ Agentless OS monitoring packages
  - Agent Builder
    - ★ Support 100+ connections for remote monitoring
    - Browser for logfile/script monitors and added CIM provider
- TADDM exploitation of ITM Agents for Discovery
- Infrastructure
  - Event Slot Customization – Improved EIF Event Integration
  - Automate agent fail-back
  - ★ Agent Management Services
  - ★ Support 64 bit counters
  - ★ Remote Deploy Bundle Performance Improvements
    - SPB bundles for TCM/TPM
    - CLI Install and Configuration of SSMs/ASMs
    - 64-bit AIX support
    - Reduced Agent Disk Footprint
    - Added situation groups
    - Enabled long situation and situation group names
  - ★ TDW schema publication tool
  - ★ FIPS 140 Compliance for TEPS
- CLI – approx. 32 new commands added, including
  - CLI for historical data configuration collection
  - Remotely invoke pdcollect tool
  - Expand tacmd createsit (display item, consecutive samples, state)
- ★ zSeries
  - Support TDW on z
  - 64-bit zLinux
  - OTEA – Support TEC Events from z HUB



Auto-generated for each key metric, per collection agent, based on learned behavior

Can adapt to important time windows based on business calendar overrides

Select situation override dialog box:

Selected schedule: Prime14to16

Calendars:

- <No schedule>
- NonPrimeShift
- Prime10to12
- Prime12to14
- Prime14to16**
- Prime16to18
- Prime8to10
- PrimeShift
- Weekday
- Weekend

Formula overrides

Expressions Details

Hostname	CPUSysTime
'thesaint'	> 68
<input checked="" type="checkbox"/> 'thesaint'	> 70
'thesaint'	> 68
'thesaint'	> 66
'thesaint'	> 40
'thesaint'	> 65

Automatic Thresholds set but provides override and centralized coordination for remote collector learning behavior

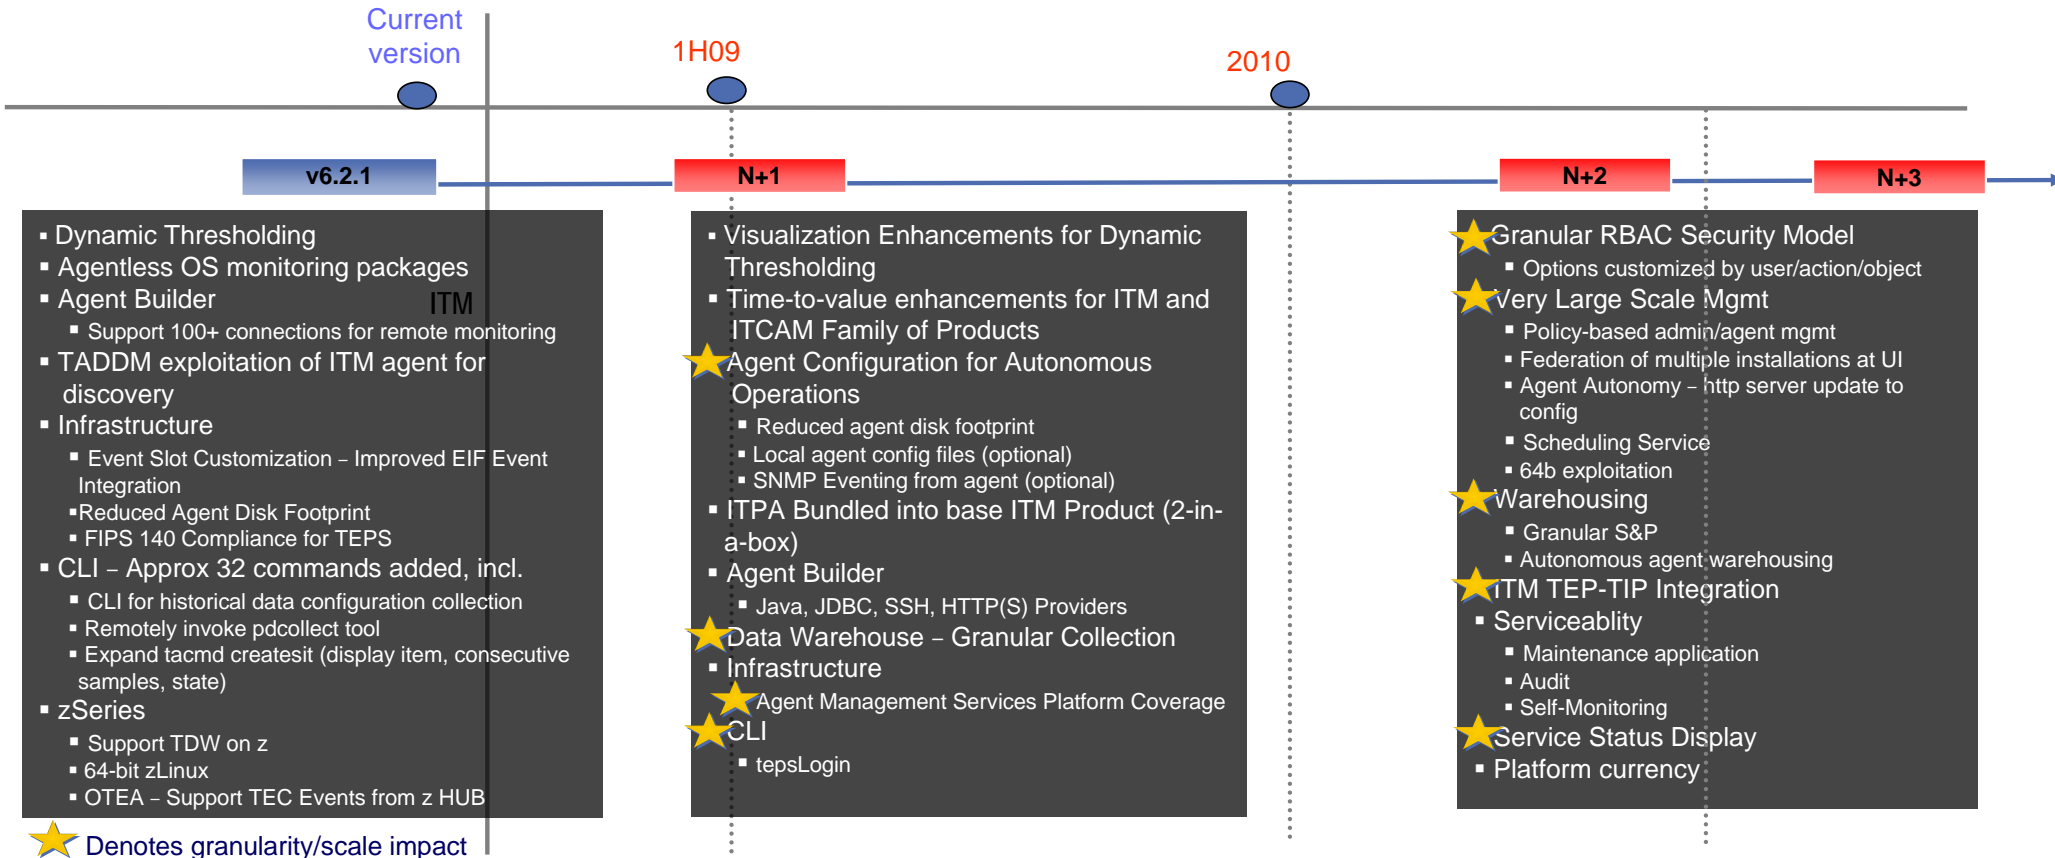
Dynamic Thresholding settings can be automatic or change controlled

# Granularity of Control

- Adaptive Monitoring – Predictive Analytics
  - Primary use case was to drive the effectiveness of monitoring:
    - Baselineing
    - Overrides
  - Secondary benefit was reduced in ITM management efforts
- Situation Groups and Managed System Lists
  - When combined, sets of heterogeneous resources can be managed by situation groups
  - Simple type matching of situations and resource types occurs dramatically improving management efficiency
- Security
  - Role-based Access Control
  - Integration with underlying security



# ITM Roadmap



★ Denotes granularity/scale impact

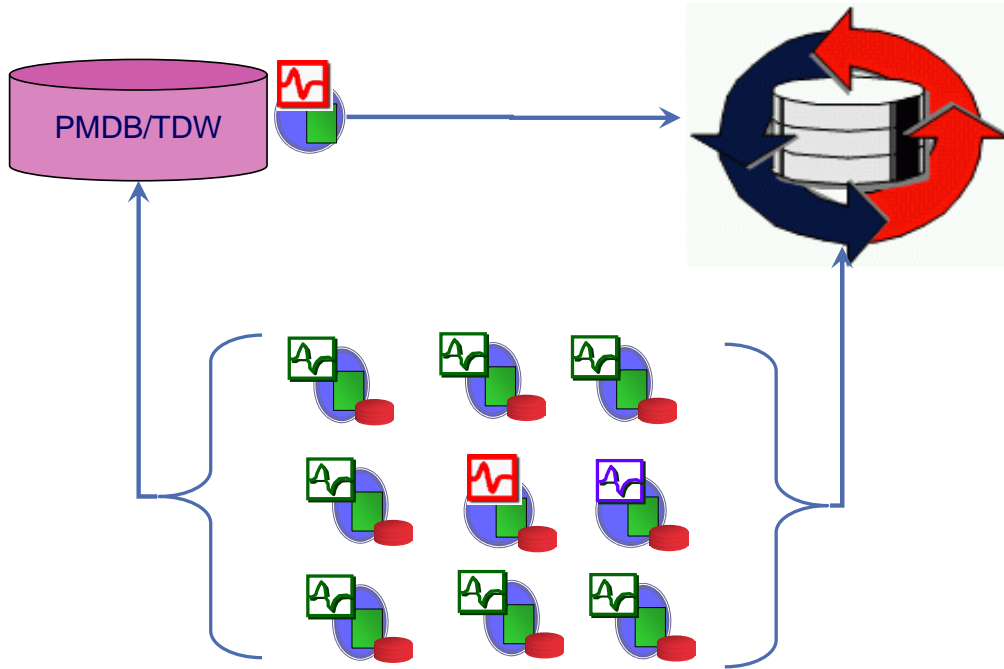


# Autonomous Agent Operations

- Customizable levels of agent behavior – embeddable (static config, exceptions-only) through integrated, centralized operations
- Capabilities
  - The ability of the agent to operate independently of a centralized infrastructure
    - Robust across intermittent and extended disconnections
      - Event data cached persistently until delivered
    - Operational configuration independently of a centralized TEMS
      - Local, “human readable” configuration files
  - Enhancements being made so they apply to a majority of agents
  - Local Configuration Options:
    - Situations – Events as SNMP traps or EIF events from the agent
    - Overrides
    - Warehousing
  - Agents may pull their most up-to-date configuration from http sources



# Predictive Analytic Agents in an OMNIBus environment



- Autonomy allows the deployment of ITM agents without a managing server and ITM UI
- ITM agents act as OMNIBus probes
- Warehousing and Reporting functions are available



# Granularity in Historical data management

- Different use cases drive different needs for historical data
  - Autonomous monitoring: local history, less than 1 month's data
  - Capacity management: centralized history, up to a few years' data
  - Predictive monitoring at service scope: centralized history
- Different business needs drive different sampling rates
  - Critical servers: Server & Process data at 1 minute intervals
  - Typical servers: Server & Process data at 15 minute intervals
  - Infrastructure: Server data at 30 minute intervals



# Granular Warehousing

The image displays two overlapping screenshots of the 'History Collection Configuration' dialog box. The top screenshot shows the 'Basic' tab with the following configuration:

- Name:** KP5\_KP509PAGIN
- Description:** Default
- Configuration:**
  - Collection Interval: 15 minutes
  - Collection Location: TEMS
  - Warehouse Interval: 1 day

The bottom screenshot shows the 'Distribution' tab with the following configuration:

- Distribute to:**  Managing System (TEMS)  Managed System (Agent)
- Start collection on:**
- Available Systems:** HUB\_SPIEGELG
- Available Managed SystemLists:** (Empty)

Both screenshots show a tree view of monitored applications on the left, including Agentless AIX OS, Agentless HP-UX OS, Agentless Linux OS, Agentless Solaris OS, Agentless Windows OS, AIX Base, AIX Premium, CCC Logs, CEC Base, Citrix Access Suite, DB2, HMC Base, i5/OS, ITM 5.x: Active Directory domain, and ITM 5.x: Active Directory replica.

- Named collection rules
- Distribution individually or by group
- Distribution to Managed System List or individual agent



# Summary – Resulting Opportunities

## Today

- Leverage Predictive Analytic Capabilities
  - Reduce time to set resource-specific baselines – Self-Learning
  - Reduce unnecessary events
    - Resource specific learning and adaptive thresholds
    - Reduction in situation maintenance times
  - Predictive Alerts with ITPA – Available now, bundled with next release
- Agent Builder Subnodes
- Infrastructure Scale Enhancements
  - Managed System Lists/Situation Groups

## Near-Term

- Performance and Event Management Synergies
- Granular Warehouse Collections – keep the relevant data



# Pulse

Comes to You 2009

IBM



*Managing the World's Infrastructure*

## Thank You!



© 2009 IBM Corporation

# Trademarks and disclaimers

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries./ Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce. ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office. UNIX is a registered trademark of The Open Group in the United States and other countries. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others. Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

© IBM Corporation 1994-2009. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

