



# The role of automated application mapping & discovery in delivering service availability

Jason A Smith

# Pulse2012

Meet the Experts. Optimise your infrastructure.

**May 31 – June 1**

Sheraton on the Park Hotel, Sydney




# Why Discovery & Dependency Mapping ?



Service  
Managers



Administrators



Operations

- **Disruptions and slow resolution still too common:**
  - *"I can't quickly determine root cause for outages or security exposures due to unknown or erroneous changes, because I can't correlate events with change and dependency information."*
- **Lack of business-driven agility exacerbates costs:**
  - *"I incur high cost and customer sat exposure as I can't determine my most critical problems as it relates to my line of business, because I don't understand how my business applications map within my IT infrastructure."*
- **Spiralling application-related costs:**
  - *"I can't reduce my data center footprint or control application costs from spiraling, because I don't understand my IT infrastructure in terms of complexity and redundancy."*
- **Lack of progress on critical projects:**
  - *"I can't implement virtualization, cloud environment, or disaster recovery centers, because I don't have the visibility into how applications are deployed in my IT infrastructure."*
- **Transition from Development to Operations is disjointed:**
  - *"The majority of applications that Operations is given don't fit our production environment or are riddled with issues."*
  - *"Our Development teams don't have a good understanding of what the production environment really looks like."*



# Can your organisation answer these key questions ?

- Often IT struggles to answer these key questions regarding the environments they manage:-
  - What do we have ?
  - How is it connected ?

## The impact if you can't

- This lack of visibility leads to.....
  - Failed changes
  - Poor problem resolution times
  - Inefficient data centres
  - Ineffective DR plans

# How Discovery & Mapping fits into IT Service Management

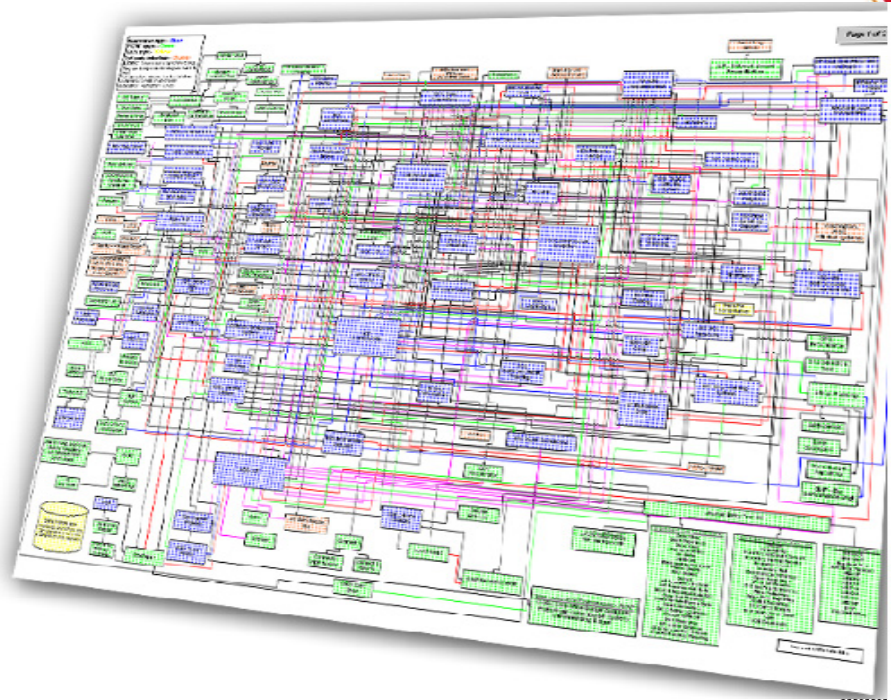
At the core of IT Service Management is visibility into the IT Services and the underlying infrastructure that supports those services.

How can you have to manage what you don't understand – you need a good understanding of what exactly makes up your services

Too often, changes result in service disruption as relationships are not immediately understood.


*Infrastructure changes = production outages*

**Visibility is key to success**

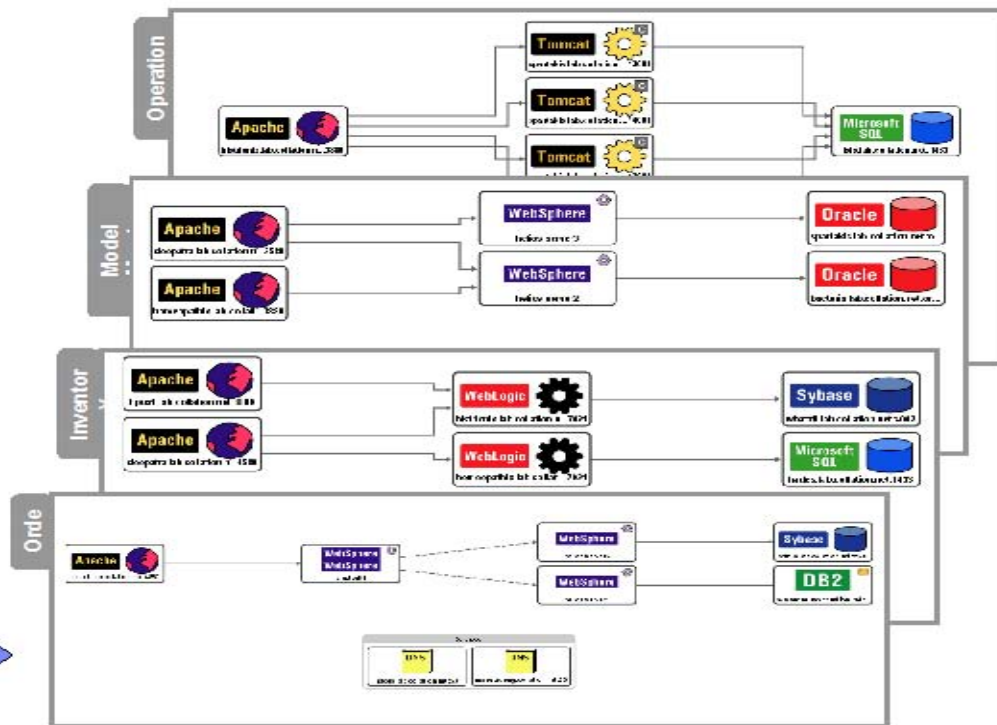


# Automated Discovery and Automated Mapping

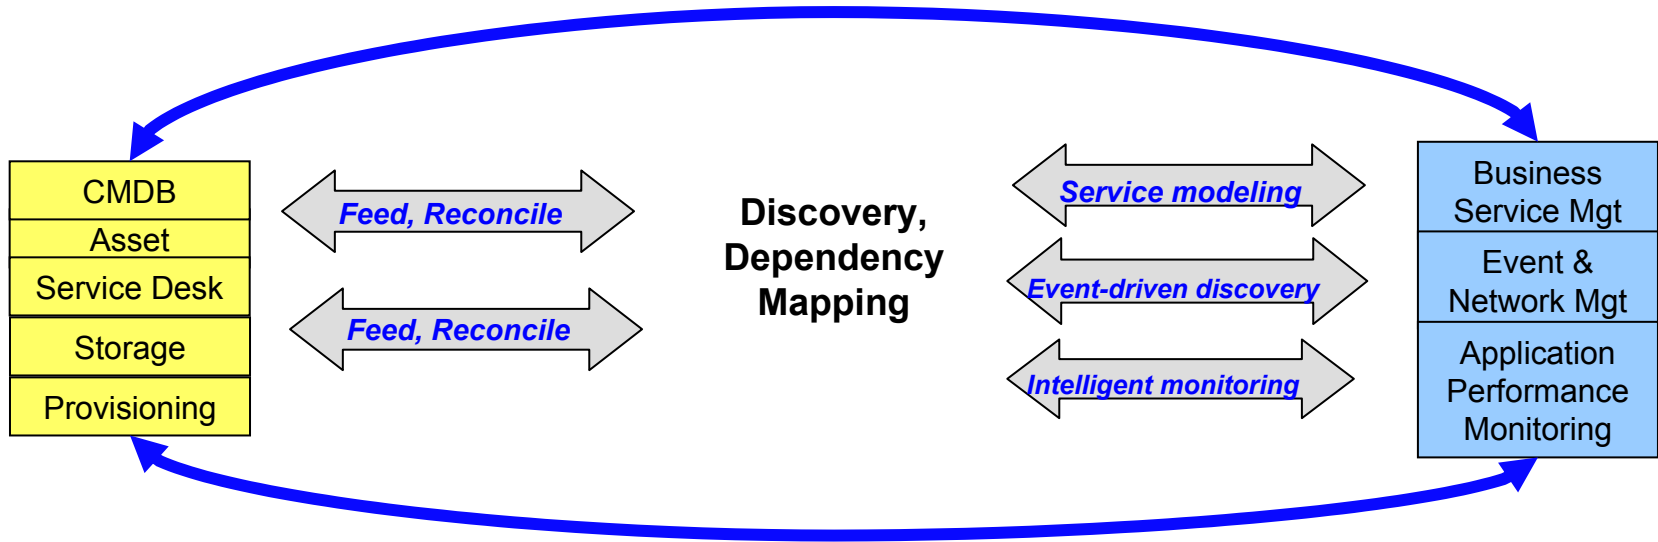
Turns this...



...into this!



# Getting started - common approaches

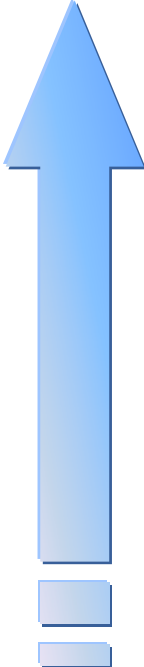


## Common customer approaches:

Configuration mgt, change mgt driven discovery  
Asset mgt reconciliation  
Storage discovery, support storage mgt tools

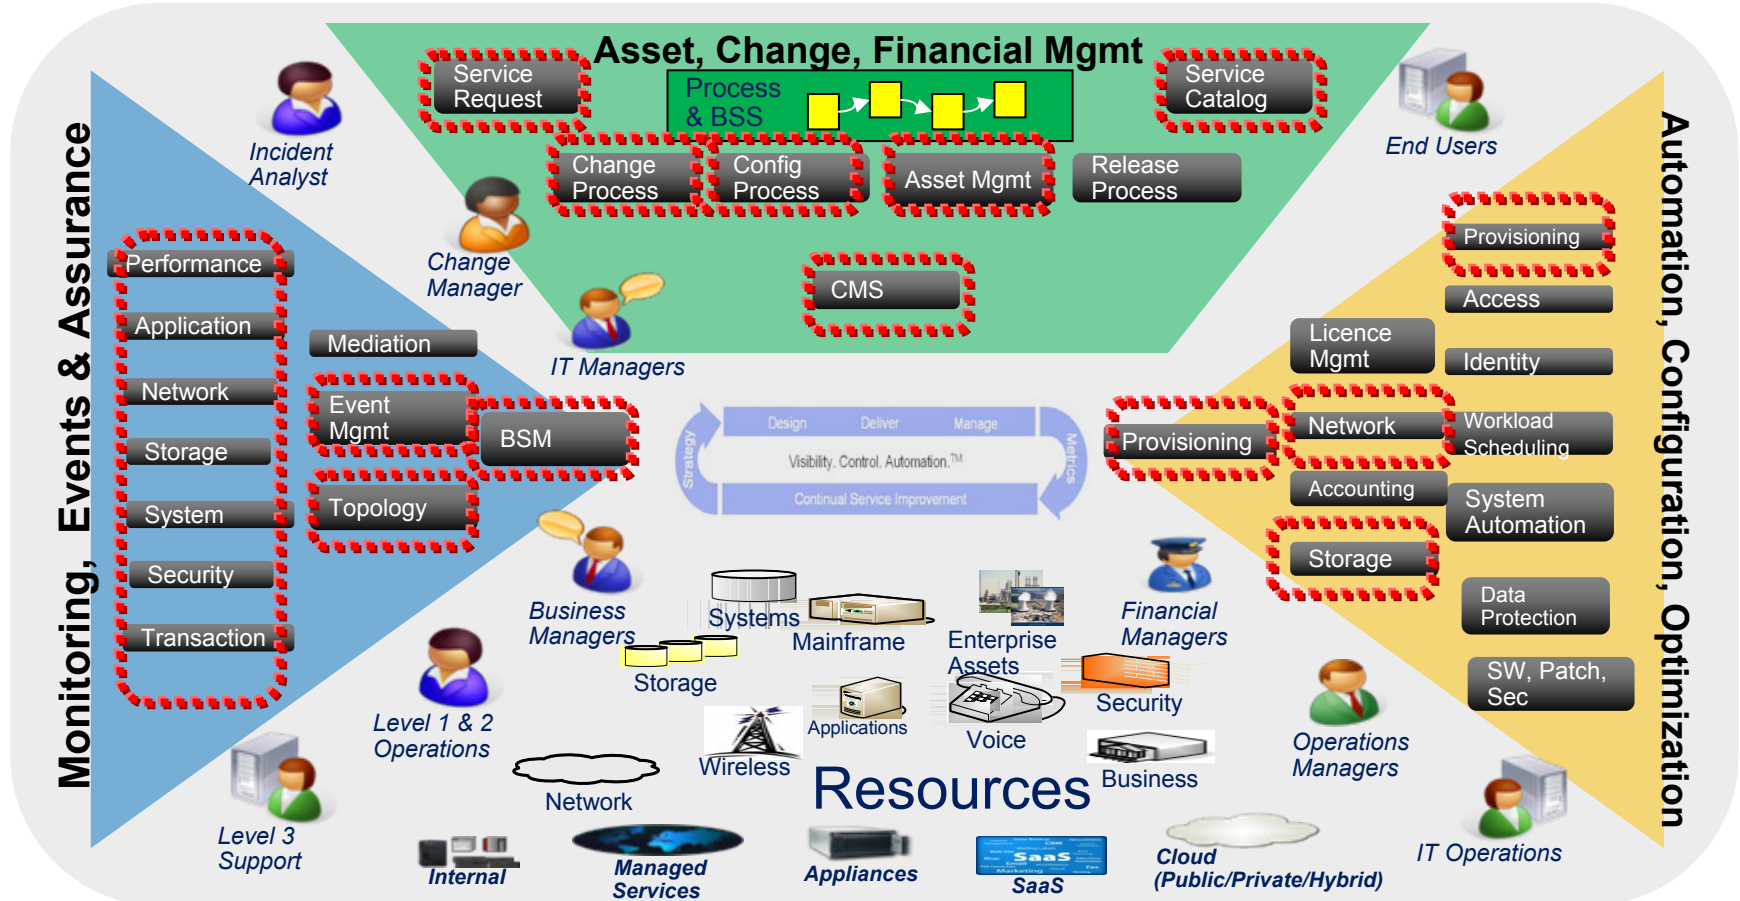
Prerequisite to Data Center Consolidation  
Intelligent monitoring and management  
Automating business service and application modeling

# Discovery is an entry point to Service Management

- 
- Start with discovery – and work your way up
    - Discover IT environment, dependencies and configurations
    - Track changes
    - Visualize dependency maps



# Where is the data used ?

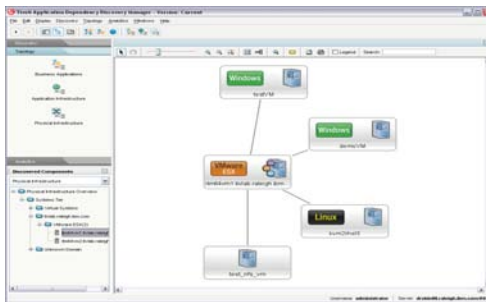


Most common areas that leverage application and infrastructure configuration and dependency data

# Tivoli Application Dependency Discovery Manager (TADDM)

## Universal Discovery Engine

**Discovers** configuration items and their Actual State. Includes Topology Views and the ability to **discover relationships** between items. **Name Reconciliation** And **Normalization** of data



## Application Mapping with Dependencies

Customer can understand what they have through agent-less **discovery of interdependencies** between applications, middleware, servers and network components and automated application maps

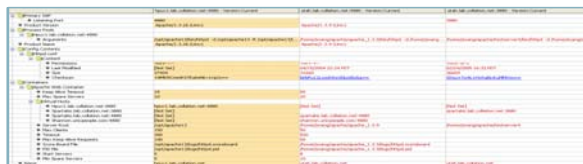
## Configuration Auditing

**Shows how configuration items are configured and changing over time** by capturing the configuration of each CI, tracking changes to it and providing analytics to report on the **history of these configuration changes** over time



## Compliance

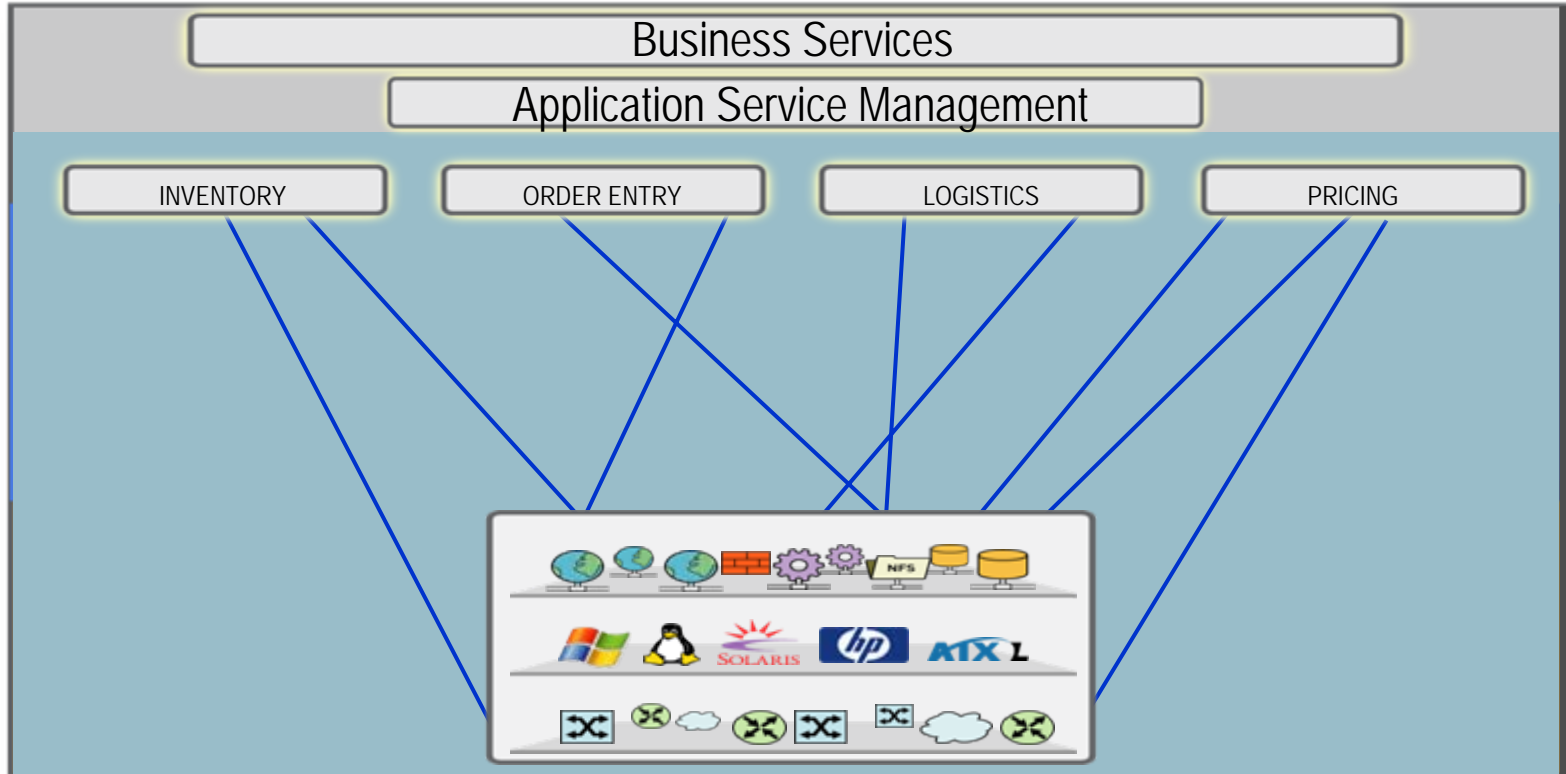
**Determines if configuration items are compliant** by using the capability to compare discovered configuration of CIs to a "reference configuration" and determine the variations that define violations to local policy

A screenshot of the TADDM software interface showing a compliance report table. The table has multiple columns, including 'Item Name', 'Status', and 'Details'. It lists various configuration items and their compliance status, with some items highlighted in yellow to indicate violations or warnings.

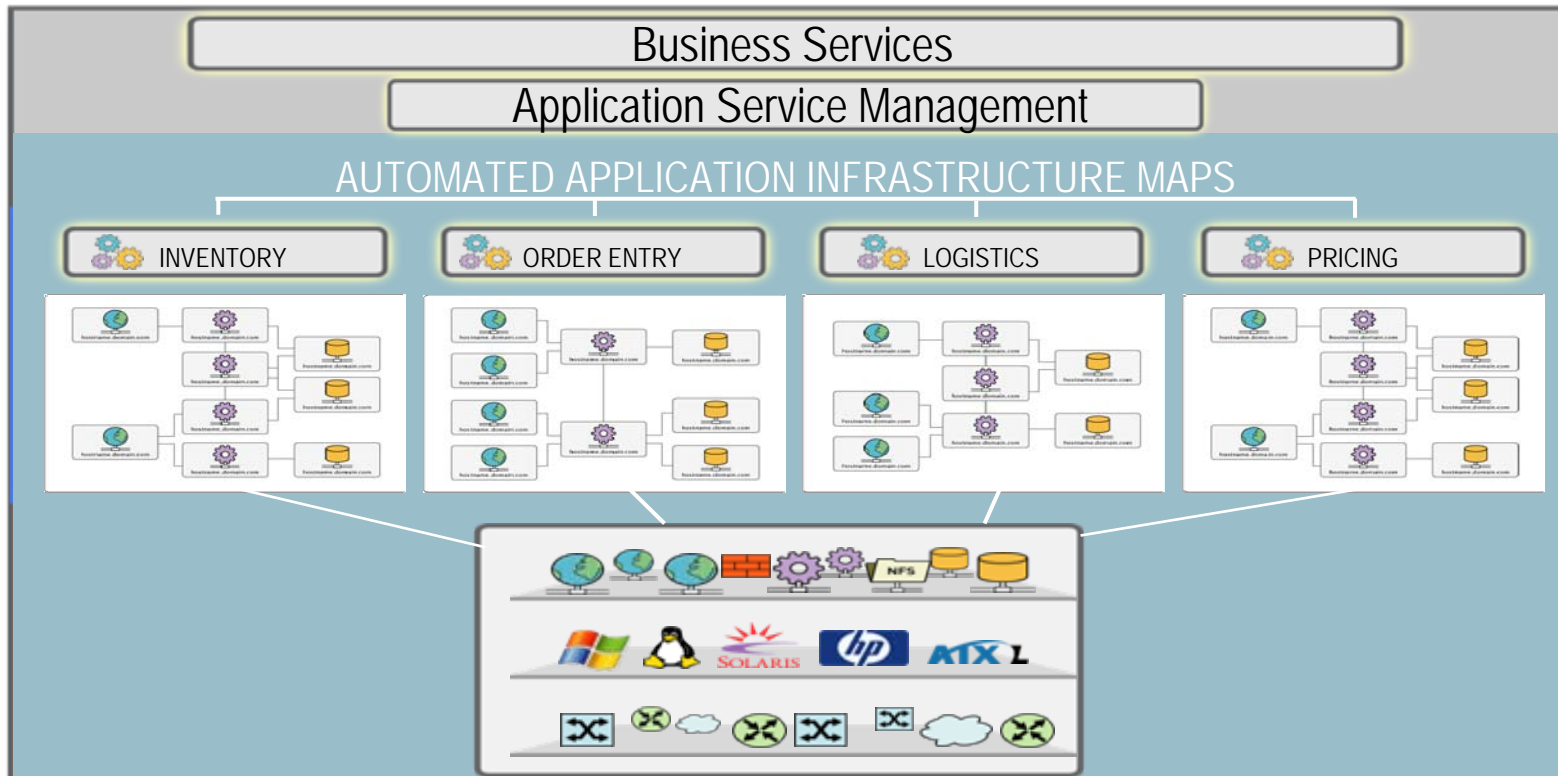
# Discovery, Dependency & Change: Integral to Operational Efficiency



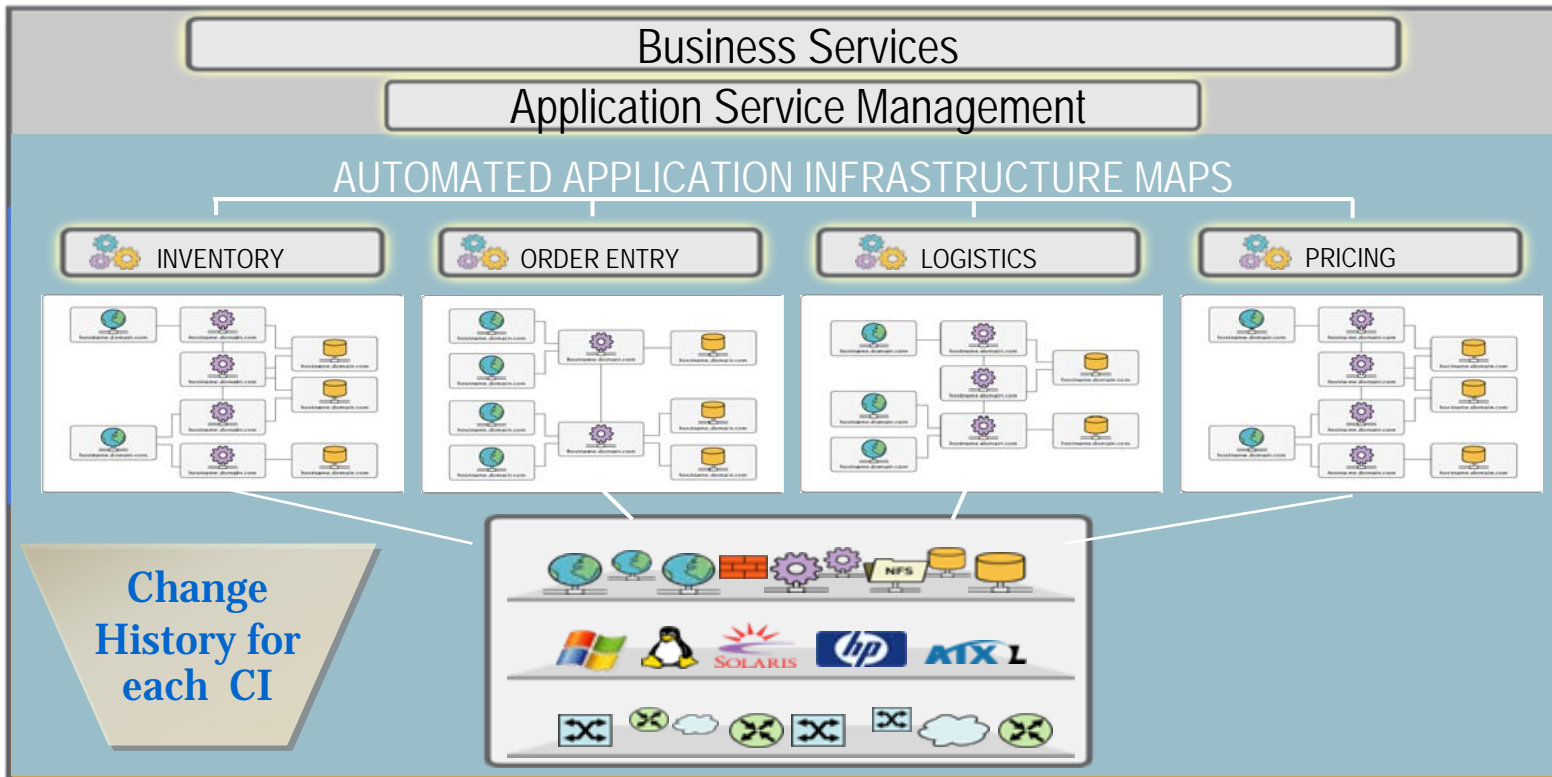
# Service Availability Challenge



# Automated Application Maps Provides the Visibility



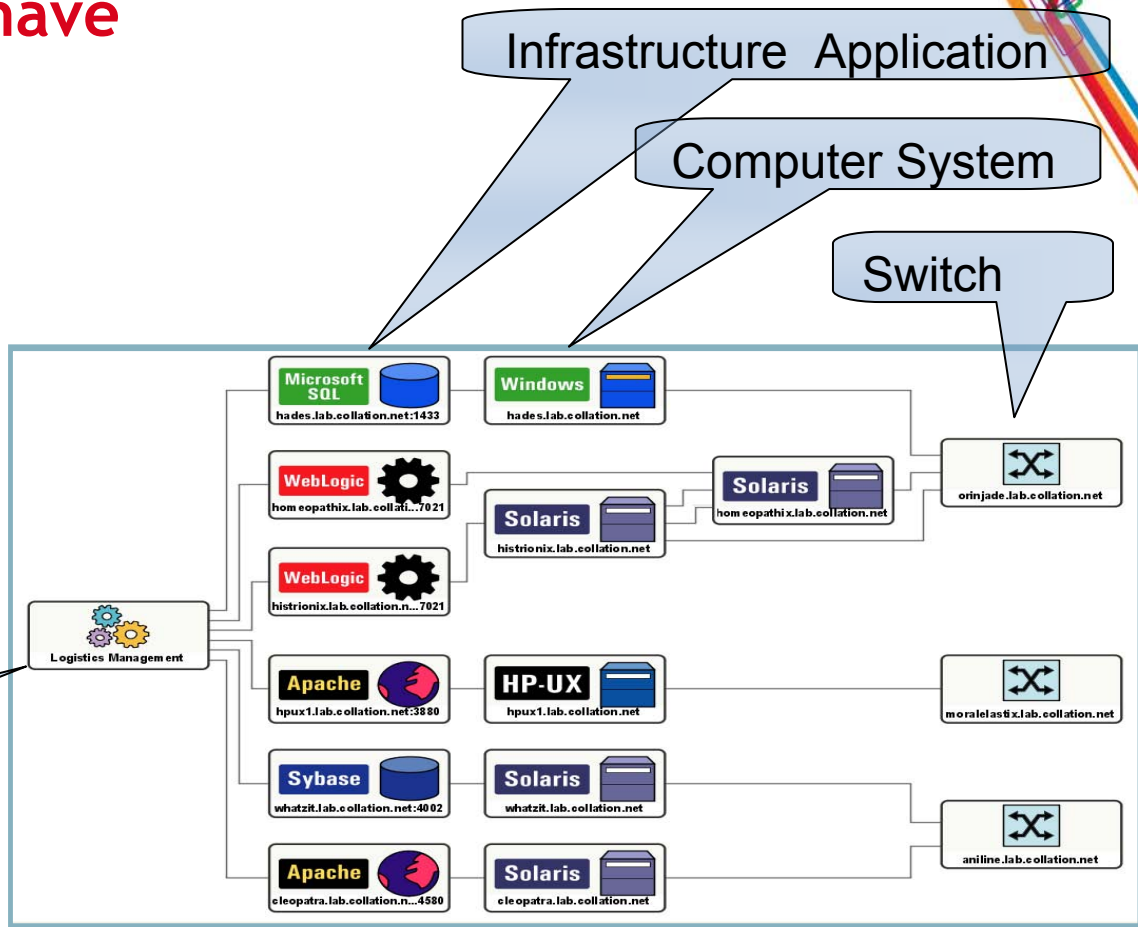
# Configuration Auditing provides the changes



# Understand what you have

- Application Mapping with Dependencies
  - Agent-less and Credential-free
  - Discover interdependencies between Applications, middleware, servers and network components)

Business Application



# Learn how your CIs are configured

- Configuration Auditing
  - Tracks changes in applications
  - Depicts that information on the map
  - Depicts that information thru reports

Automatically tracks changes on all CIs & attribute values over time...

Application

Type ▾	Component	Change	Date	Attribute	Old Value	New Value
Apache	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	appDescriptors		/usr/local/apache/appd
Apache	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	appDescriptors		/usr/local/apache//app
ApacheWebContainer	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	ApacheWebContainer	/usr/local/apache/	/usr/local/apache
ApacheWebContainer	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	ApacheWebContainer	15	20
ApacheWebContainer	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	ApacheWebContainer	88	100
ProcessPool	homeopathix.lab.collati	Updated	12/04/2004 15:01 PST	homeopathix.lab.collati	/usr/local/apache//bin/	./httpd -d /usr/local/a



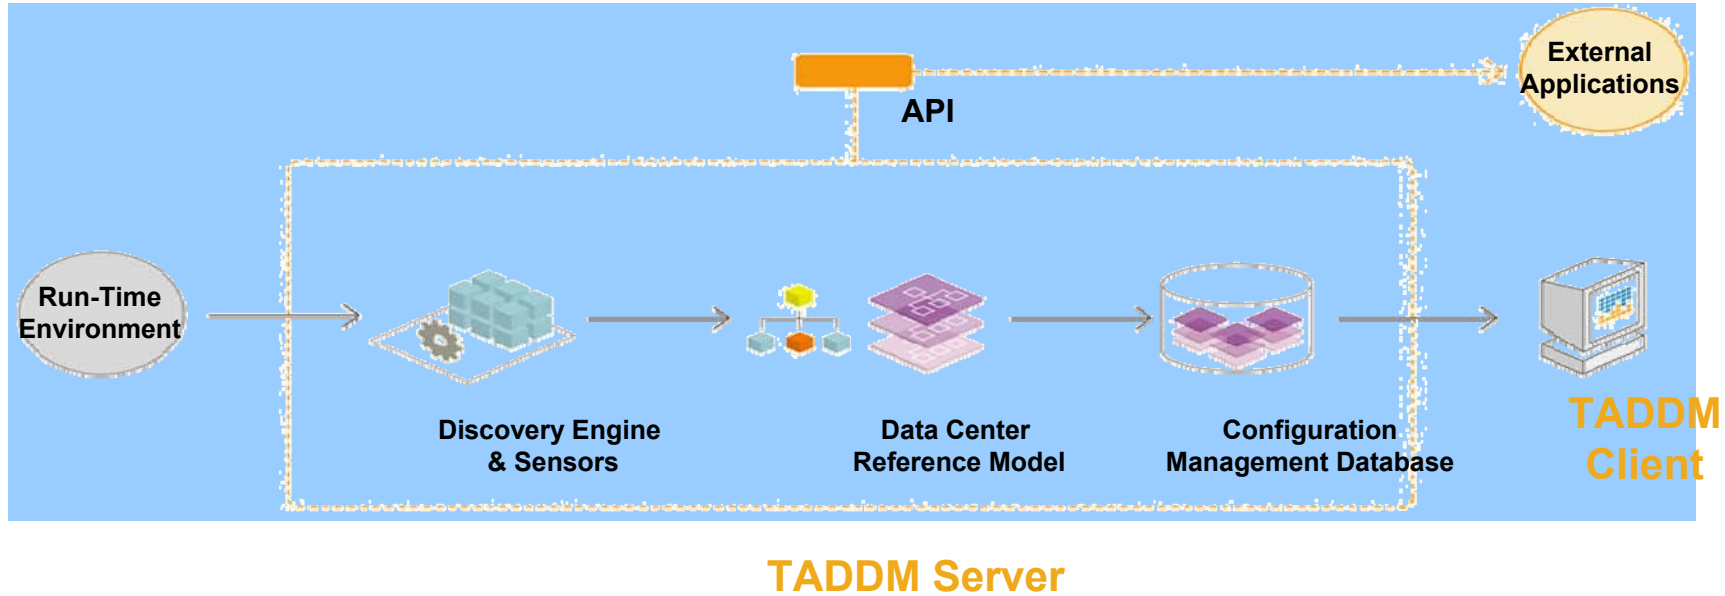
# Determine if it is compliant

- Compliance
  - Compare configuration to “reference master”
  - Compare to your standard policy

Comparing two instances of an Apache Web Server to the reference master

	hpux1.lab.collation.net:4880 - Version:Current	utah.lab.collation.net:4880 - Version:Current	utah.lab.collation.net:3880 - Version:Current
Primary SAP			
Listening Port	4880	4880	3880
Product Version	Apache/1.3.26 (Unix)	Apache/1.3.9 (Unix)	Apache/1.3.9 (Unix)
Process Pools			
Hpx:1.lab.collation.net:4880			
Arguments	/opt/apache13/bin/httpd -d /opt/apache13 -R /opt/apache13/l...	/home/jwang/apache/apache_1.3.9/bin/httpd -d /home/jwang...	/home/jwang/apache/testserver4/bin/httpd -d /home/jwang/a...
Product Name	Apache/1.3.26 (Unix)	Apache/1.3.9 (Unix)	Apache/1.3.9 (Unix)
Config Contents			
Httd.conf			
Content			
Permissions	-rwxr-----	-rwxr-xr-x	-rwxr-xr-x
Last Modified	[Not Set]	04/15/2004 22:24 PDT	02/24/2006 16:03 PDT
Checksum			
Containers			
Apache Web Container			
Keep Alive Timeout	15	55	55
Max Spare Servers	10	20	20
Virtual Hosts			
Hpx:1.lab.collation.net:4880	hpux1.lab.collation.net:4880	[Not Set]	[Not Set]
Spartakis.lab.collation.net:3880	[Not Set]	spartakis.lab.collation.net:4880	spartakis.lab.collation.net:3880
Shannon.unixpeople.com:4880	[Not Set]	shannon.unixpeople.com:4880	shannon.unixpeople.com:4880
Server Root	/opt/apache13	/home/jwang/apache/apache_1.3.9	/home/jwang/apache/testserver4
Max Clients	150	50	50
Timeout	300	500	500
Max Keep Alive Requests	100	50	50
Score Board File	/opt/apache13/logs/httpd.scoreboard	/home/jwang/apache/apache_1.3.9/logs/httpd.scoreboard	/home/jwang/apache/apache_1.3.9/logs/httpd.scoreboard
PID file	/opt/apache13/logs/httpd.pid	/home/jwang/apache/apache_1.3.9/logs/httpd.pid	/home/jwang/apache/apache_1.3.9/logs/httpd.pid
Start Servers	5	10	10
Min Spare Servers	5	10	10
Name	hpux1.lab.collation.net	utah.lab.collation.net	utah.lab.collation.net

# How TADDM Works



# TADDM Features

## ➤ Performance and scalability:

- New streaming architecture

## ➤ New discovery capabilities:

- Script-based discovery
- Asynchronous discovery
- Concurrent targeted discovery

## ➤ New business application grouping composer:


- Simplifies business application definition and discovery collection
- Based on user-specified criteria

## ➤ Enhanced interoperability:


- Improved TBSM Integration;
- Proactive discovery
- TWS discovery scheduling

## Enhanced UI and Reporting:


Quickly create custom reports



## Credential-leverage discovery



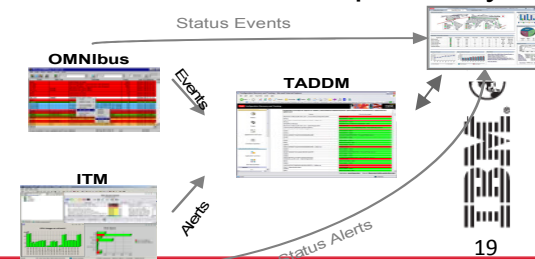
## Grouping Composer



## Concurrent discovery



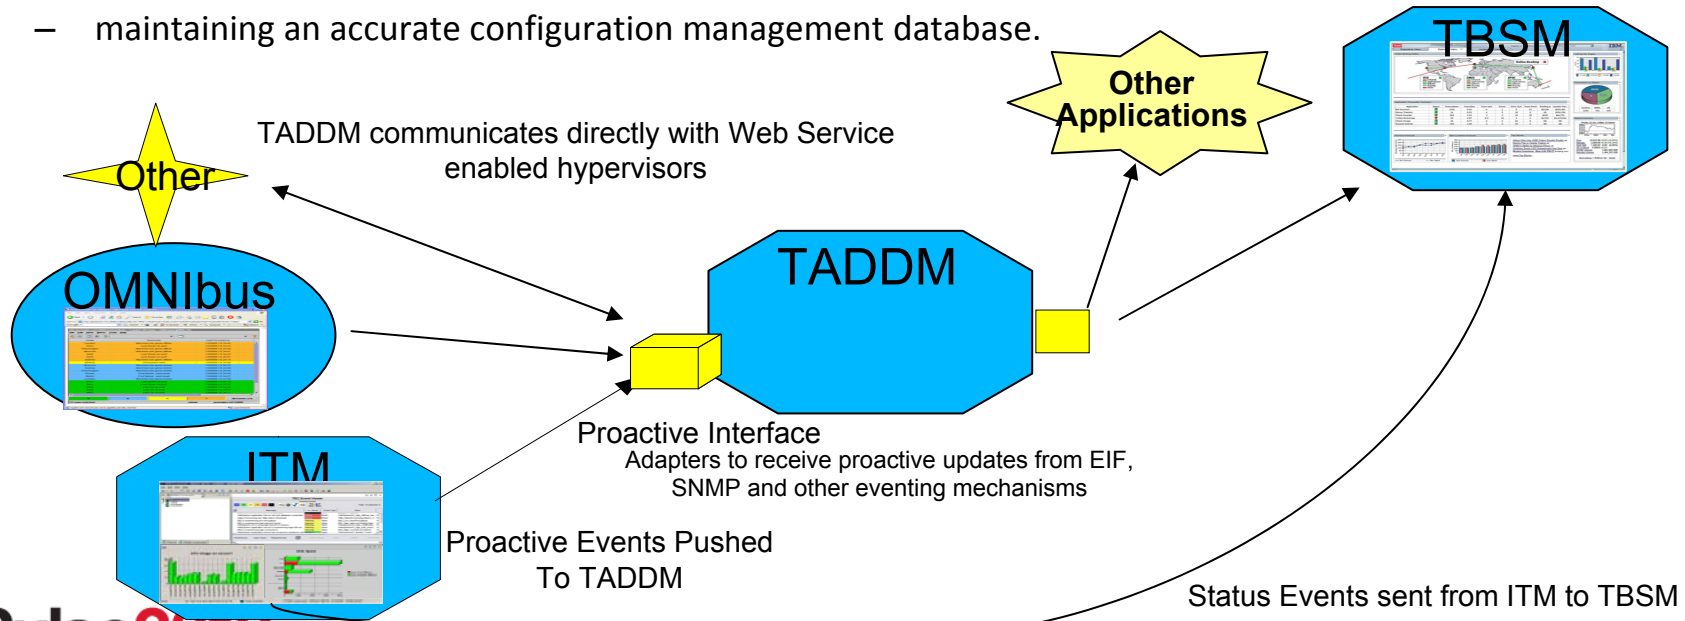
## Enhanced Interoperability



# Proactive Discovery


Provides accurate, near real-time visibility into the datacenter from an application standpoint. Critical for:

- planning data center consolidations
- managing configuration change impact
- virtualization initiatives
- maintaining an accurate configuration management database.



# Concurrent Discovery

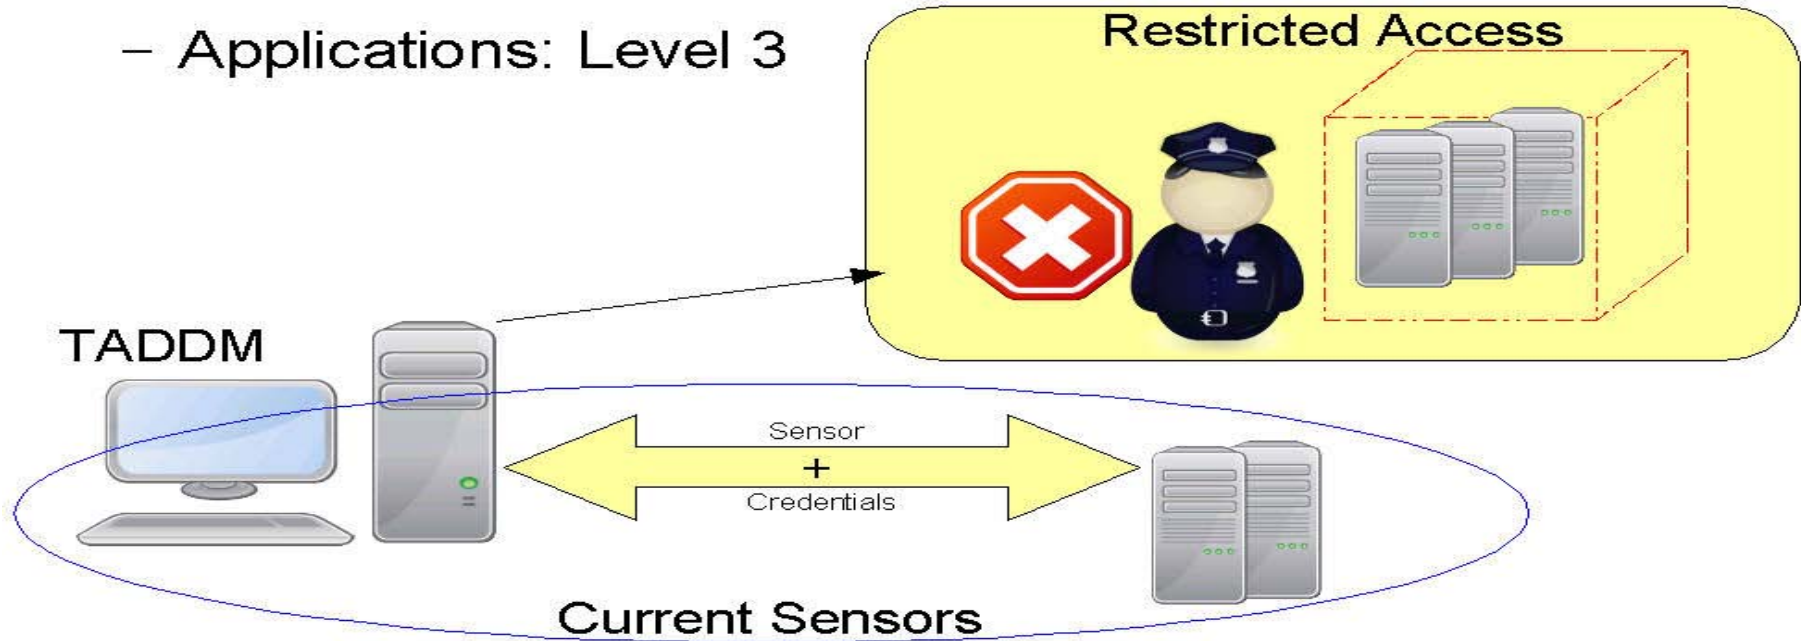
- Run multiple discovery jobs supported on a single discovery server.
- Allows for targeted small scope refreshes in parallel to scheduled long running discoveries.



*Benefits: Improved discovery efficiency and scalability  
Added flexibility to run targeted discovery runs as needed*

# Common Challenge: Credentials are restricted, difficult to obtain

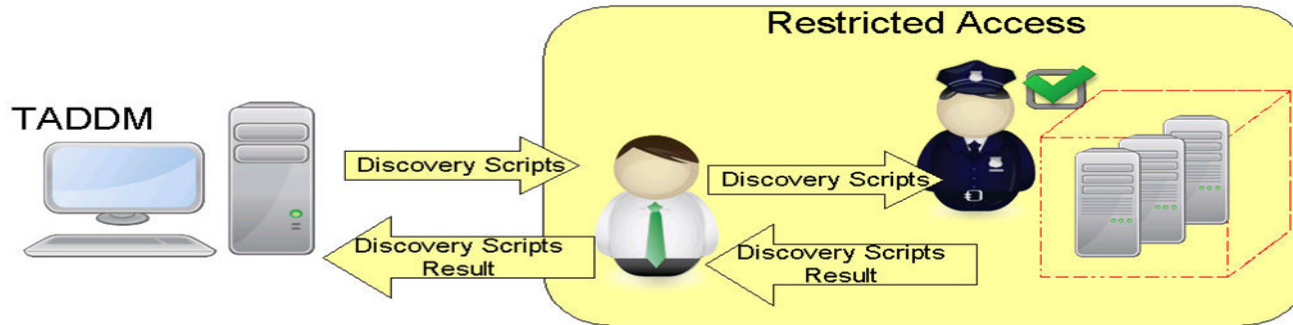
- Current Sensors -> Credentials:
  - Computer System: Level 2
  - Applications: Level 3



## Asynchronous Discovery

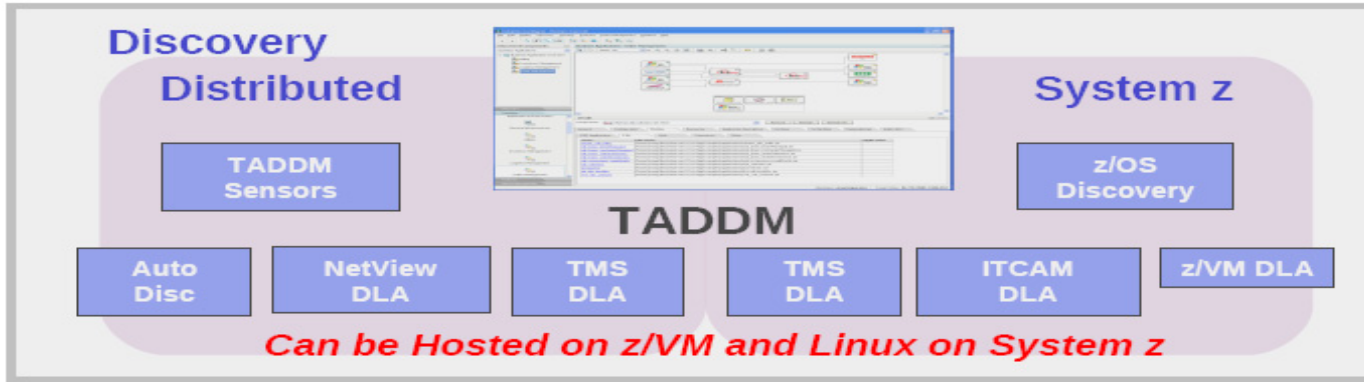
Provides flexibility to run discovery scans leveraging existing middleware credentials and access control parameters

- Retrieve the script(s) to run from TADDM server
- Put results in TADDM server



## Automated Discovery of z and Distributed Resources

- Create / discover / visualize Business Applications and their dependencies between z/OS, z/VM and distributed resources
- Maintain and track configuration changes
- Baseline gold standard and measure configuration drift
- Comparison of configurations (test, user, QA, production)



**Tivoli Application Dependency Discovery Manager dynamically gathers configuration information from both Distributed and z/OS Resources and map their relationships.**



## Trademarks and disclaimers

© Copyright IBM Australia Limited 2012 ABN 79 000 024 733 © Copyright IBM Corporation 2012 All Rights Reserved. TRADEMARKS: IBM, the IBM logos, ibm.com, Smarter Planet and the planet icon are trademarks of IBM Corp registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

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.

