IBM Research

Research Report

Hari Ramasamy

IBM T.J. Watson Research Center

IFIP 10.4 WG Meeting

9/19/08

Two Parts

- **Pointers to our recent research results towards realizing the vision of tomorrow’s data centers**
- **Ongoing efforts in *automatic discovery technologies* and leveraging them for *availability management***

IBM Research

Customer Isolation in Multi-Tenant Data Centers

Hari Ramasamy

Joint Work with

Matthias Schunter, Bernhard Jansen, Konrad Eriksson (IBM Zurich Research Lab)

Serdar Cabuk, Chris Dalton (HP Labs)

IFIP 10.4 WG Meeting 9/19/08

Today’s Data Centers

- **Most Fortune 500 companies have their own large, dedicated data centers**
- **Smaller companies are increasingly outsourcing their IT infr., but still “physical cages” model at the data center provider**
- **Over-provisioning, over-engineering, under-utilization rampant**

What might tomorrow's data center look like? [IBM NEDC White Paper]

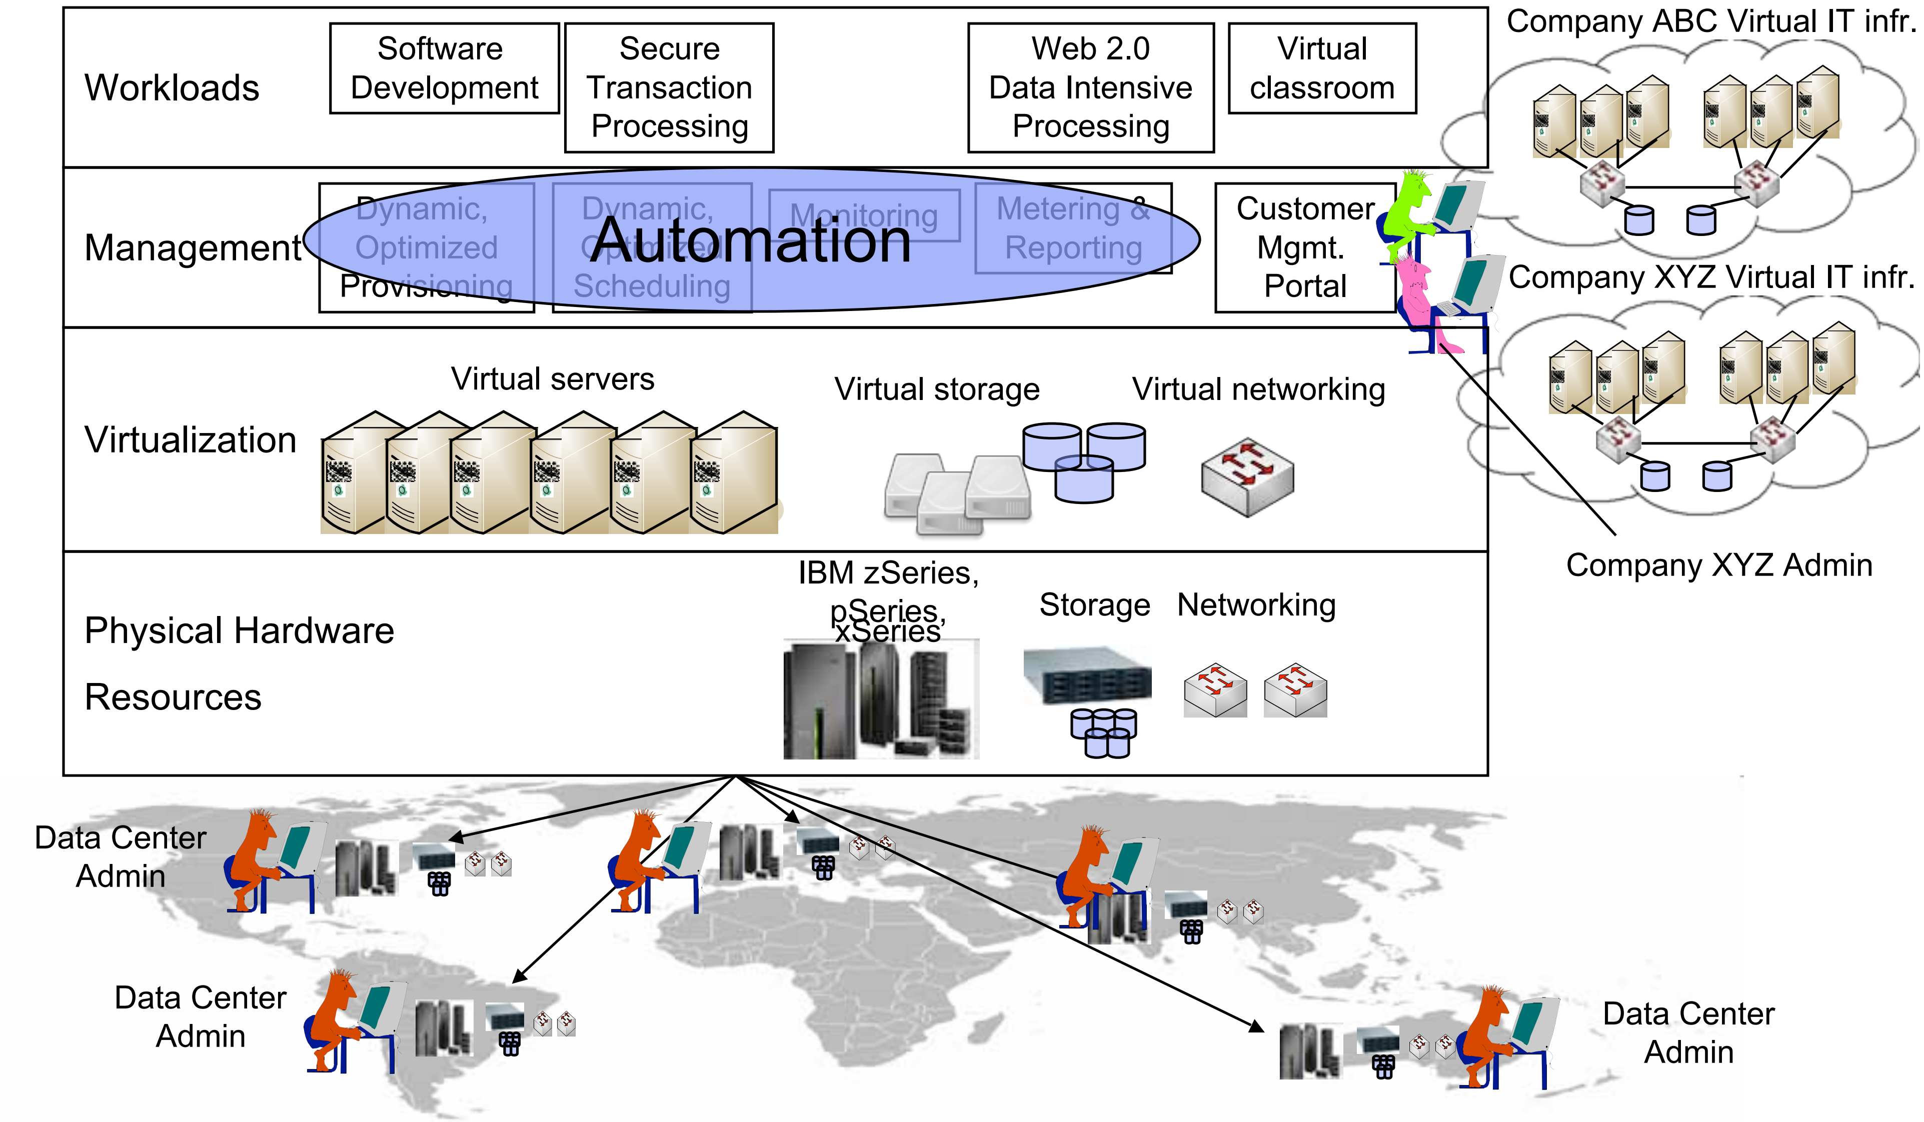

Workloads

Software Development

Secure Transaction Processing

Web 2.0 Data Intensive Processing

Virtual classroom

Management

Dynamic, Optimized Provisioning

Dynamic, Optimized Scheduling

Monitoring

Metering & Reporting

Automation

Customer Mgmt. Portal

Virtualization

Virtual servers

Virtual storage

Virtual networking

Physical Hardware

Resources

IBM zSeries, pSeries, xSeries

Storage

Networking

Company ABC Virtual IT infr.

Company XYZ Virtual IT infr.

Company XYZ Admin

Data Center Admin

Data Center Admin

Data Center Admin

Customer Isolation in Multi-Tenant Data Centers

- **Isolation policies have multiple aspects**
  - e.g., networking, storage, and VM lifecycle management
- **How do you enforce the isolation policies in a unified manner?**
  - **Trusted Virtual Domains**
- **How do you ensure the integrity of policy enforcement components and evaluate their trustworthiness?**

The Trusted Virtual Domain (TVD) Model

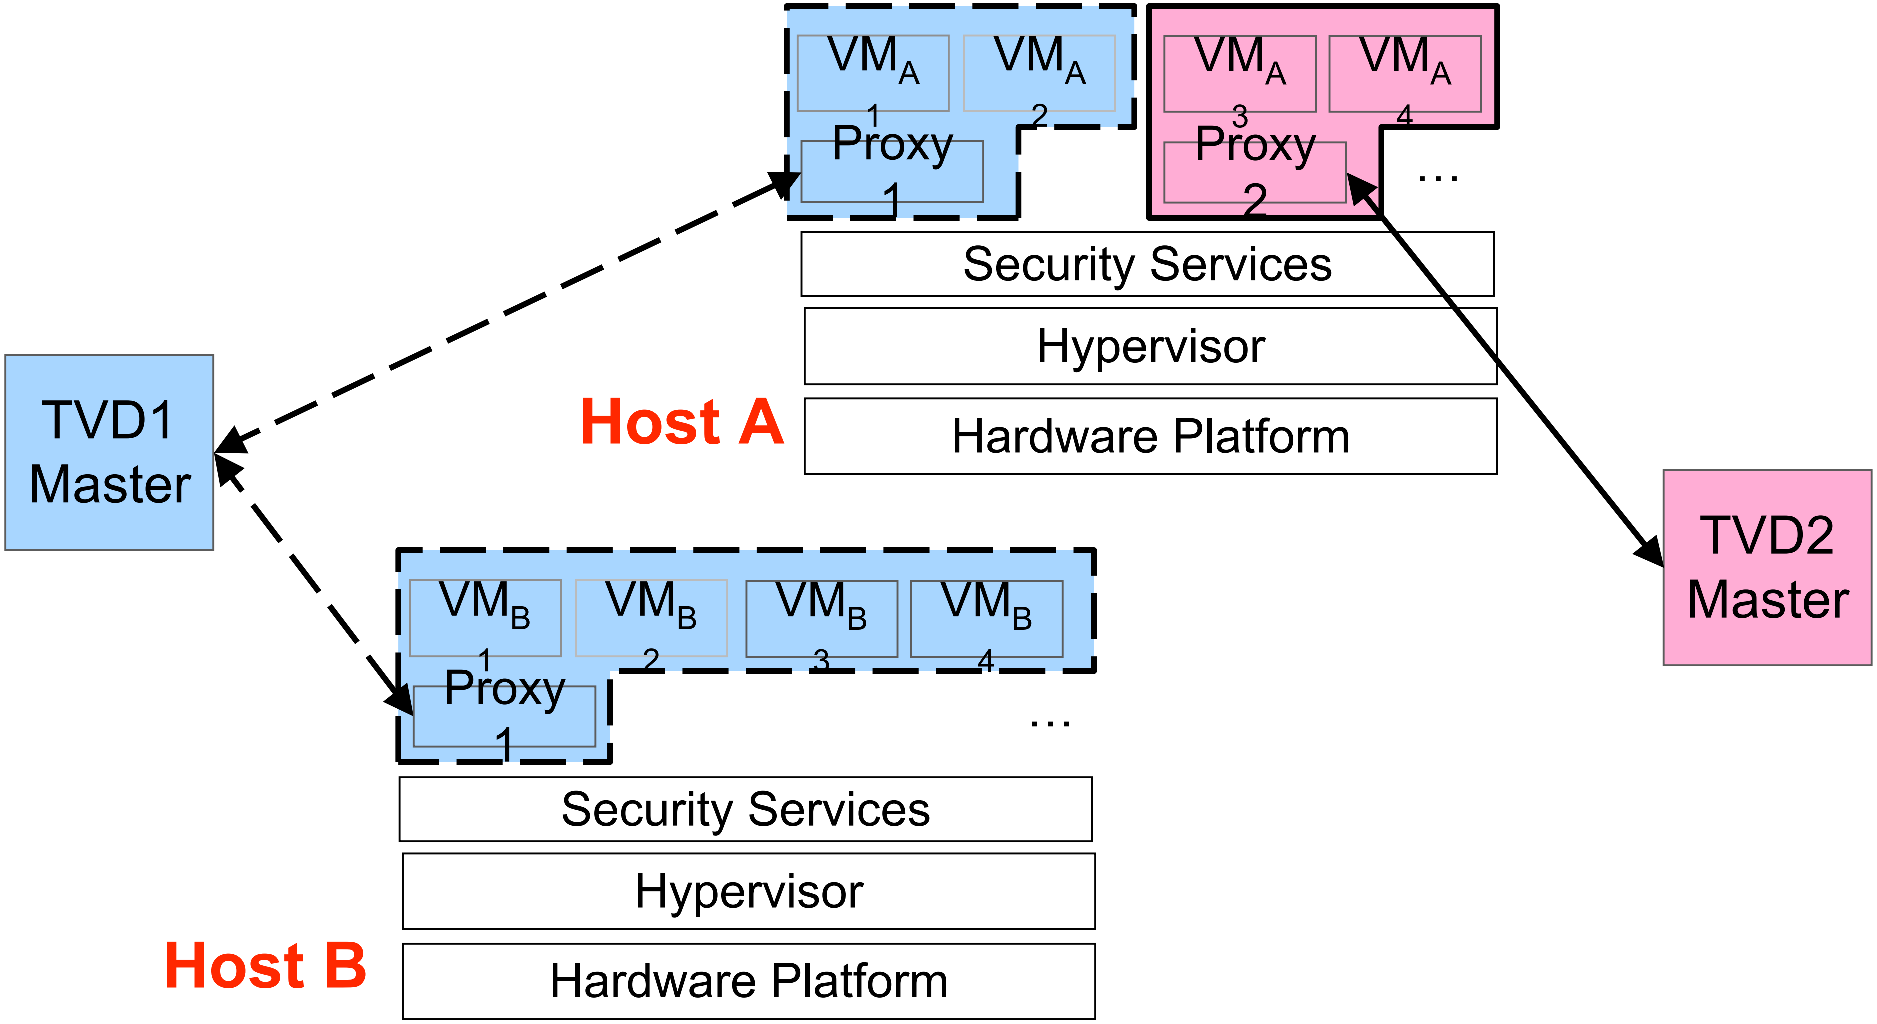

Customer Isolation in Multi-Tenant Data Centers

- **Isolation policies have multiple aspects**
  - e.g., networking, storage, and TVD membership
- **How do you enforce the isolation policies in unified manner?**
  - Trusted Virtual Domains
    - Towards Automated Provisioning of **Secure Virtualized Networks** [ACM CCS '07]
      - Map high-level isolation policies into information flow control policies
      - Automatically instantiate virtual networking elements (VPN, VLAN tagging, EtherIP encapsulation) for each VM to satisfy flow control policies
      - Xen-based prototype implementation
- **How do you ensure the integrity of policy enforcement components and evaluating their trustworthiness?**

Secure Network Virtualization: Xen-based Prototype

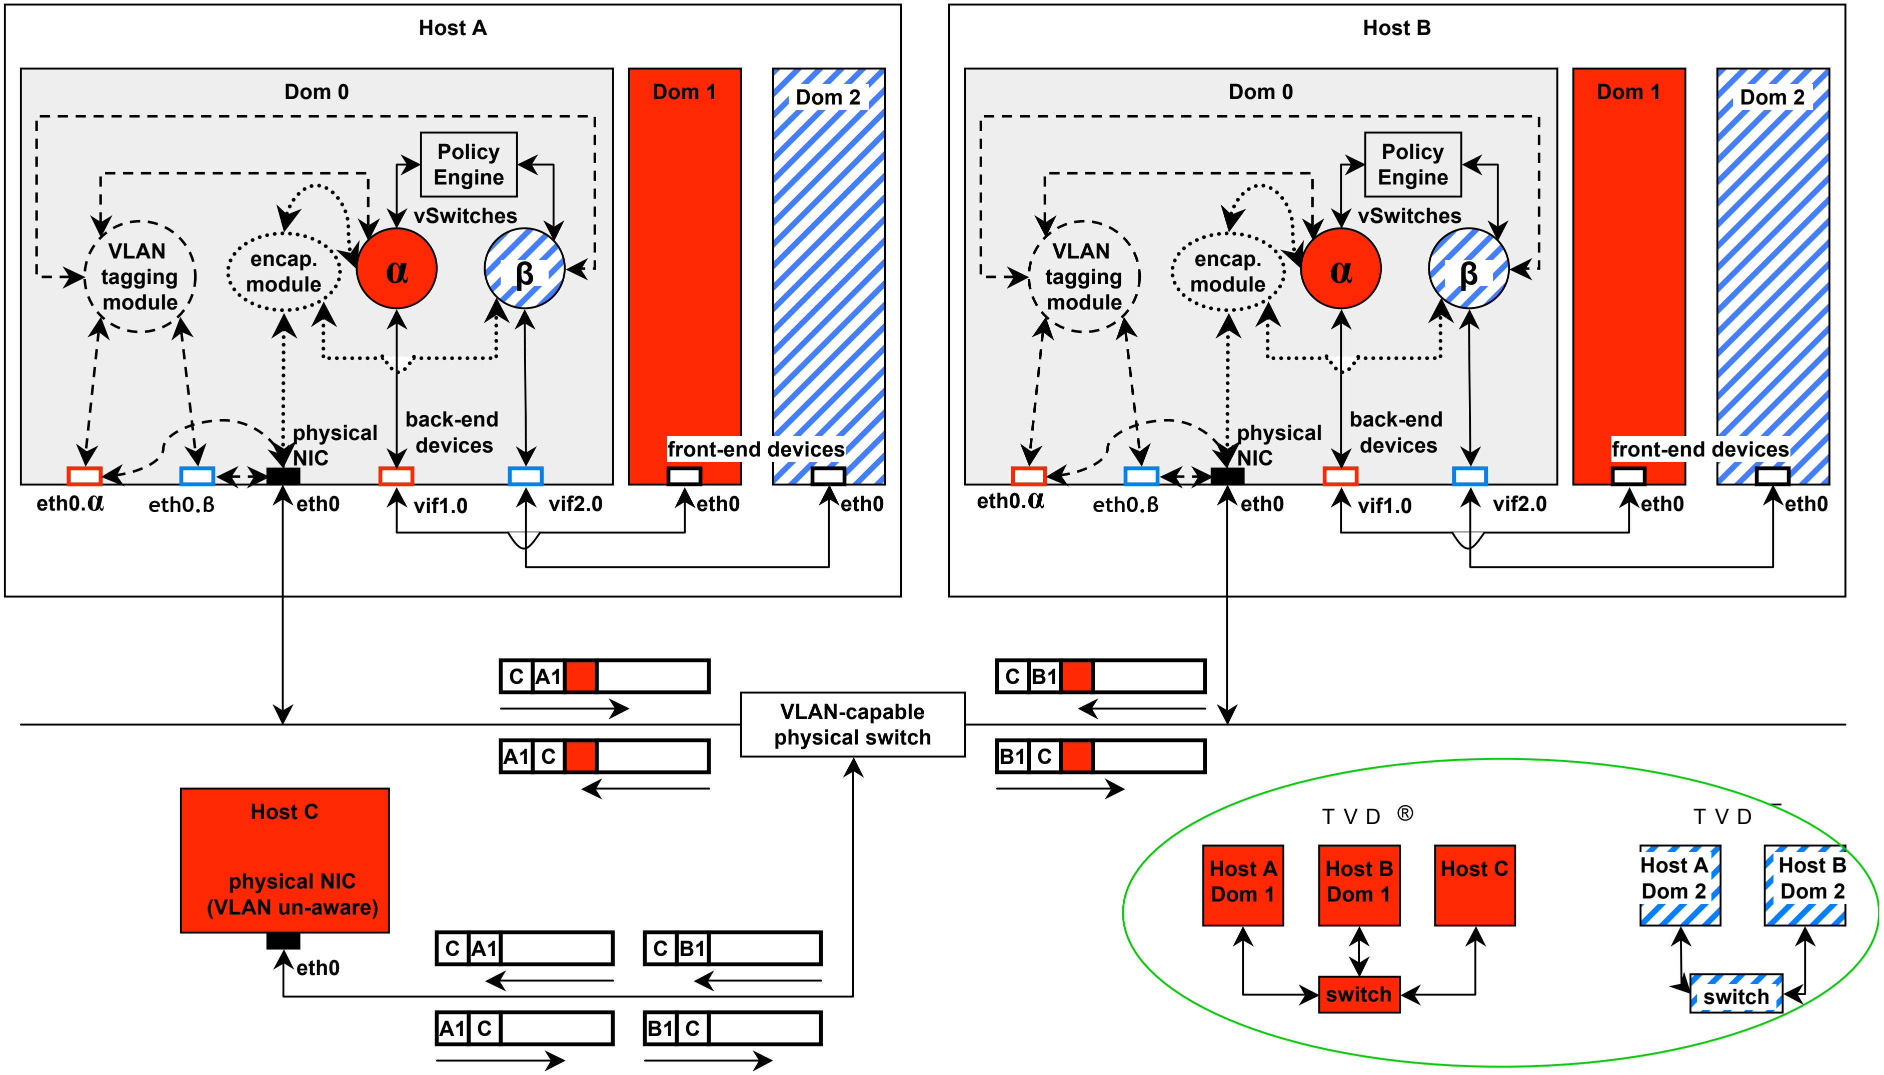

Customer Isolation in Multi-Tenant Data Centers

- **Isolation policies have multiple aspects**
  - e.g., networking, storage, and TVD membership
- **How do you enforce the isolation policies in unified manner?**
  - Trusted Virtual Domains
    - Towards Automated Provisioning of **Secure Virtualized Networks** [ACM CCS ‘07]
- **How do you ensure the integrity of policy enforcement components and evaluating their trustworthiness?**
  - Hardware-based Trusted Computing
    - Policy Enforcement and Compliance Proofs for Xen Virtual Machines [ACM VEE ‘08]
      - Fine-grained compliance checks that can be expressed as predicates on log entries
      - Xen-based prototype implementation

Two Parts

- **Pointers to our recent research results towards realizing the vision of tomorrow’s data centers**
- **Ongoing efforts in *automatic discovery technologies* and leveraging them for *availability management***

IBM Research

Using Automated Discovery for *End-to-End* Availability Management

Hari Ramasamy

Joint Work with

Murthy V. Devarakonda, Nikolai Joukov, Kostas Magoutis,

Birgit Pfitzmann, and Norbert G. Vogl

IFIP 10.4 WG Meeting

9/19/08

Introduction

- **End-to-end availability ≈ availability as the user perceives it**

- **1st step in availability mgmt: Determine "as-is" state (discovery)**
  - What are the services, apps, IT infrastructure components?
  - What are the relationships between them?

- **Why is discovery even a problem?**
  - Doesn't model-based deployment solve this problem?
  - Can't the sys admin or a consultant tell you?

Discovery Tools

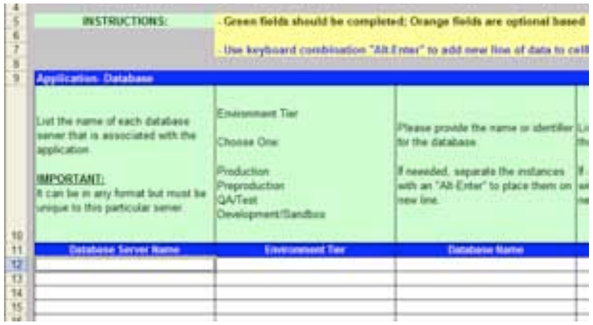

INSTRUCTIONS: - Green fields should be completed; Orange fields are optional based

- Use keyboard combination "Alt-Enter" to add new line of data to cell

Application: Database

List the name of each database server that is associated with the application. IMPORTANT: It can be in any format but must be unique to this particular server	Environment Tier Choose One: Production Preproduction QA/Test Development/Sandbox	Please provide the name or identifier for the database If needed, separate the instances with an "Alt-Enter" to place them on new line.
Database Server Name	**Environment Tier**	**Database Name**

- **Manual Discovery**
  - most common today
  - incomplete, inaccurate, quickly becomes outdated

- **Network-Level Automated Discovery**
  - network mapping (e.g., NMAP)
  - traffic analysis (e.g. analysing CISCO Netflow packets)
  - E.g., TADDM level 1, Aurora, eMulsa (all IBM), EMC ADM
- **System-Level Automated Discovery**
  - issue shell commands, instrument code, pass tags across layers
  - agent-based (resident code) or agent-less
  - E.g., TADDM levels 2 and 3, HP DDM, **Galapagos**

Example End-to-End Dependencies Using Galapagos (conceptual picture)

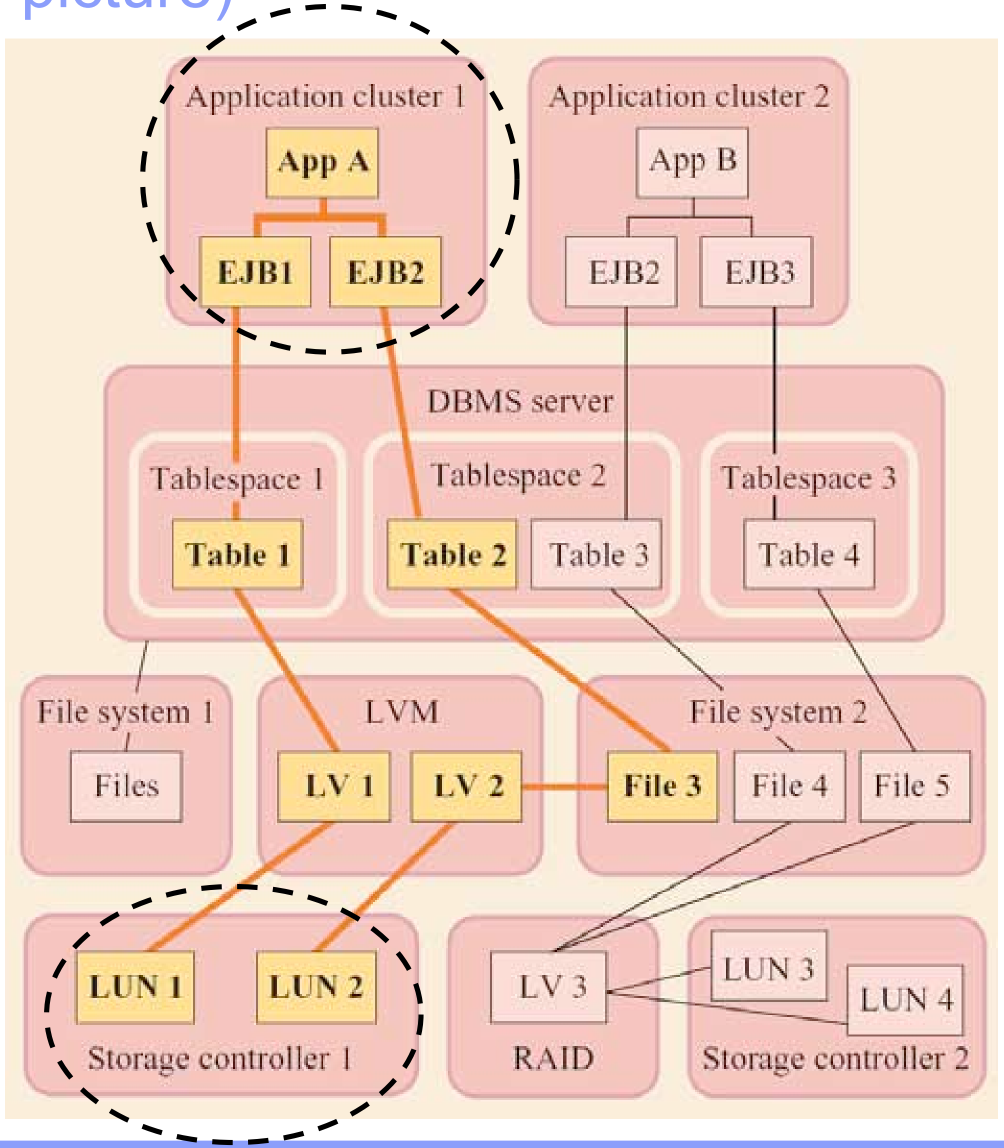

End-to-end, deep, *application-data* and *data-data* relationships

Example End-to-End Dependencies Using Galapagos (real output)

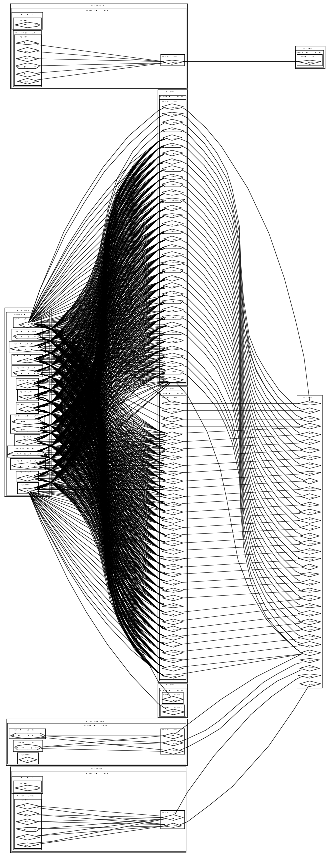

IBM HTTP Server

WAS cells, nodes, app. servers, applications, modules,

DB2 database servers, instances, databases

Galapagos Data Location Templates (DLTs)

- **Library of pre-created DLTs**
- **One DLT per software component**
  - DB2, IBM WebSphere App. Server, IBM HTTP Server, etc.
  - DLT models data-related artifacts of that component
- **DLTs are platform-independent**
  - But, refer to sensors which give installation-specific info
- **DLTs are reusable, composable**

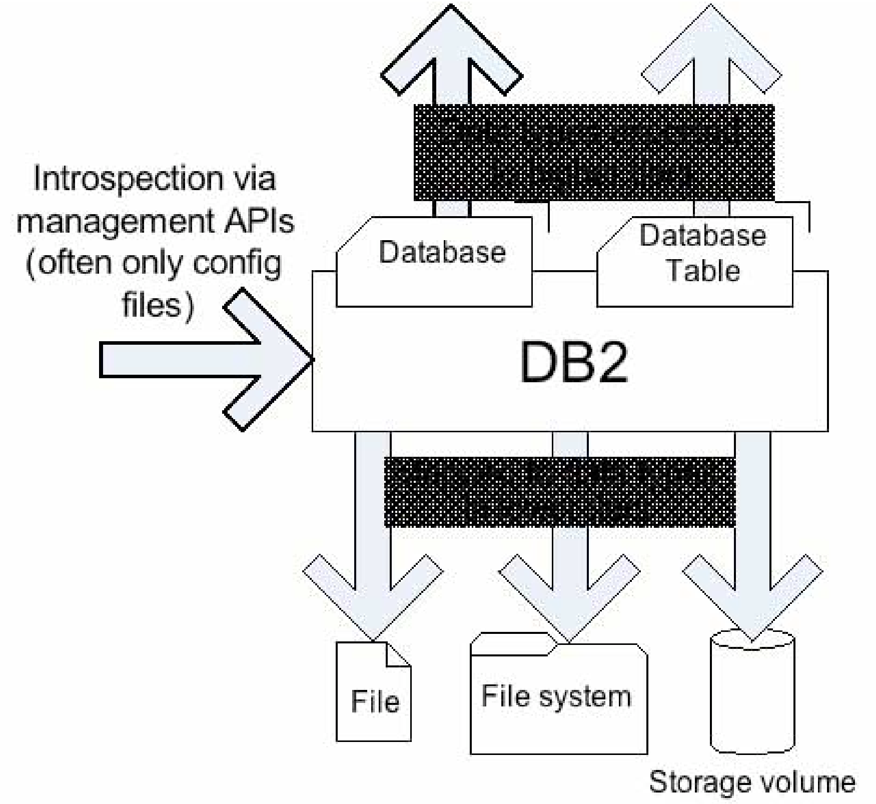

Galapagos Discovery Process [IBM J. R&D 2008]

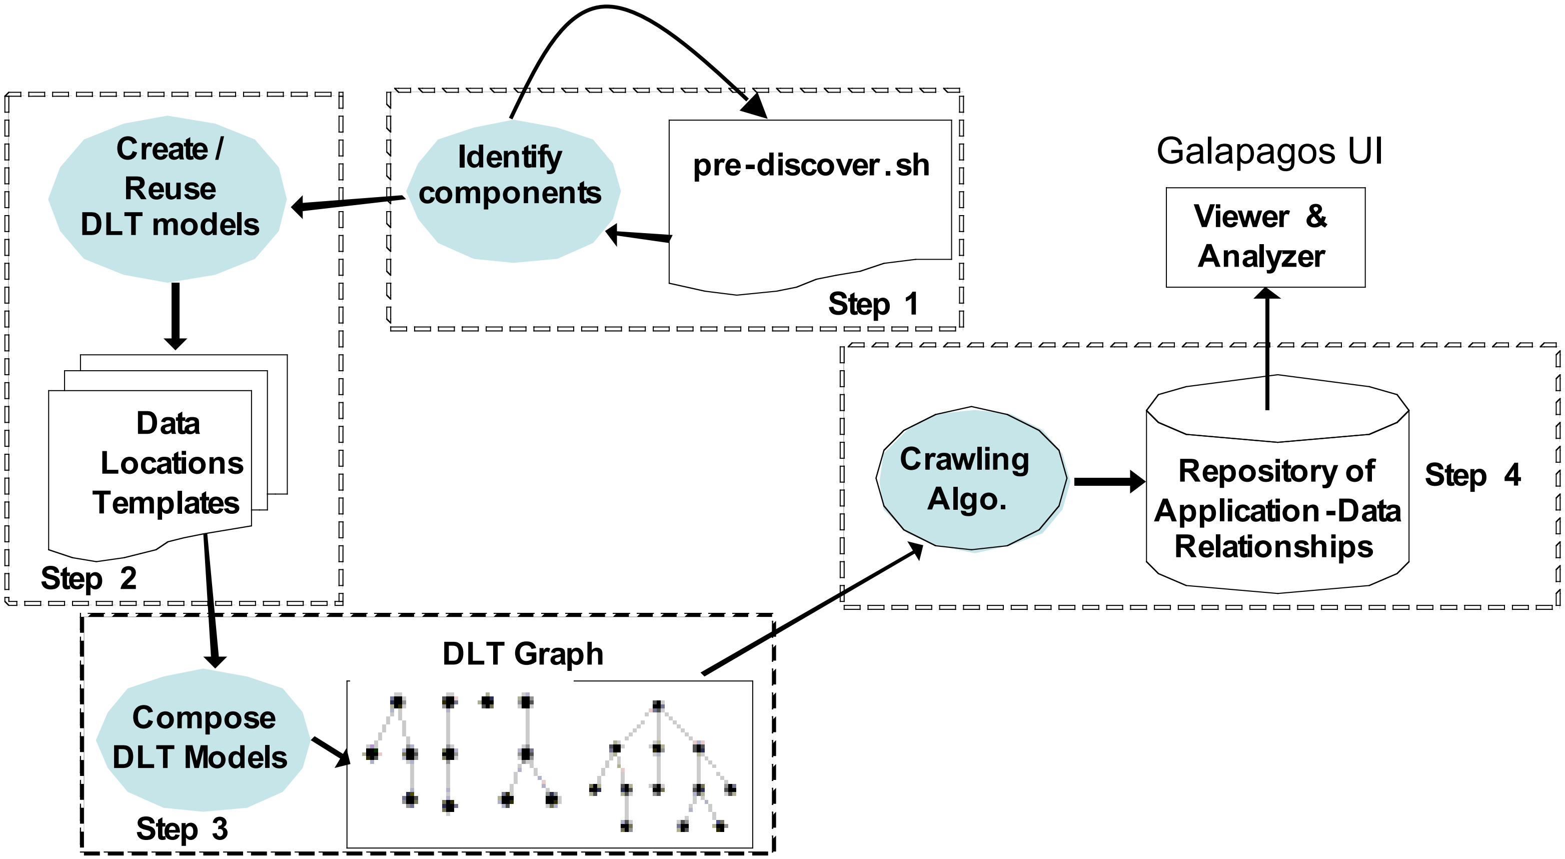

Applying Discovery for Availability Management

- **End-to-End Availability Assessment**
- **Recovery Planning**
- **Identifying Actually Enforced Failure Boundaries**
- **Re-structuring a business process into availability domains**
- **Reducing the scope of sensitivity analysis**
- **Problem Determination**

Identifying Actually Enforced Failure Boundaries

- Failure independence is a key simplifying assumption (e.g., using Reliability Block Diagrams)

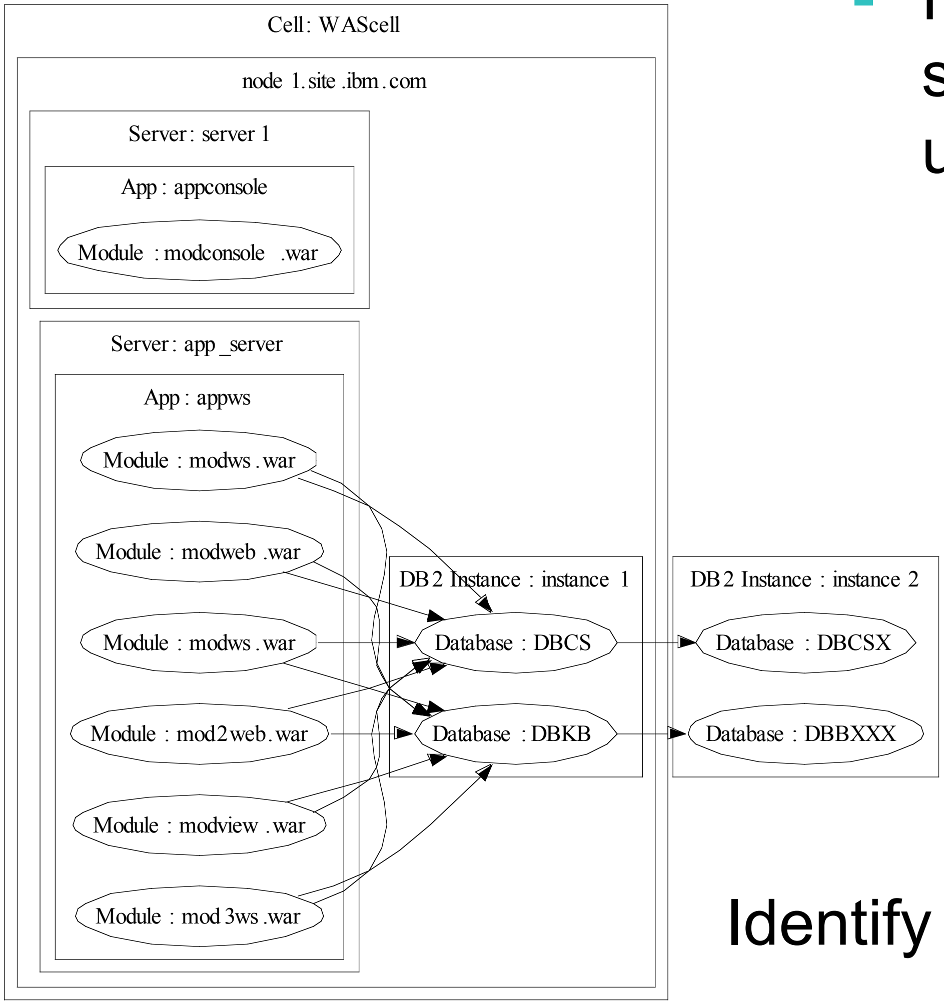

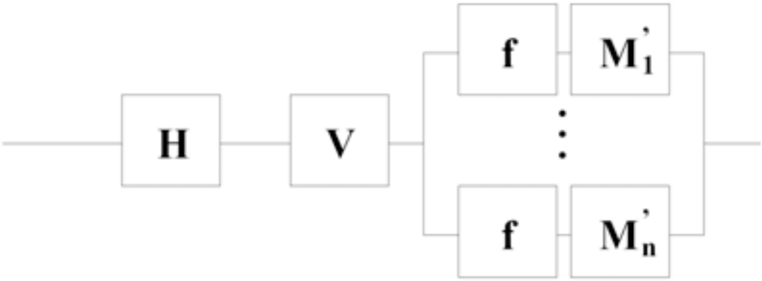

Identify co-located applications, DBs

Ongoing work on identifying co-located VMs

Recovery Planning

- **Dependency information helps determine sequence of tasks for recovering a service**

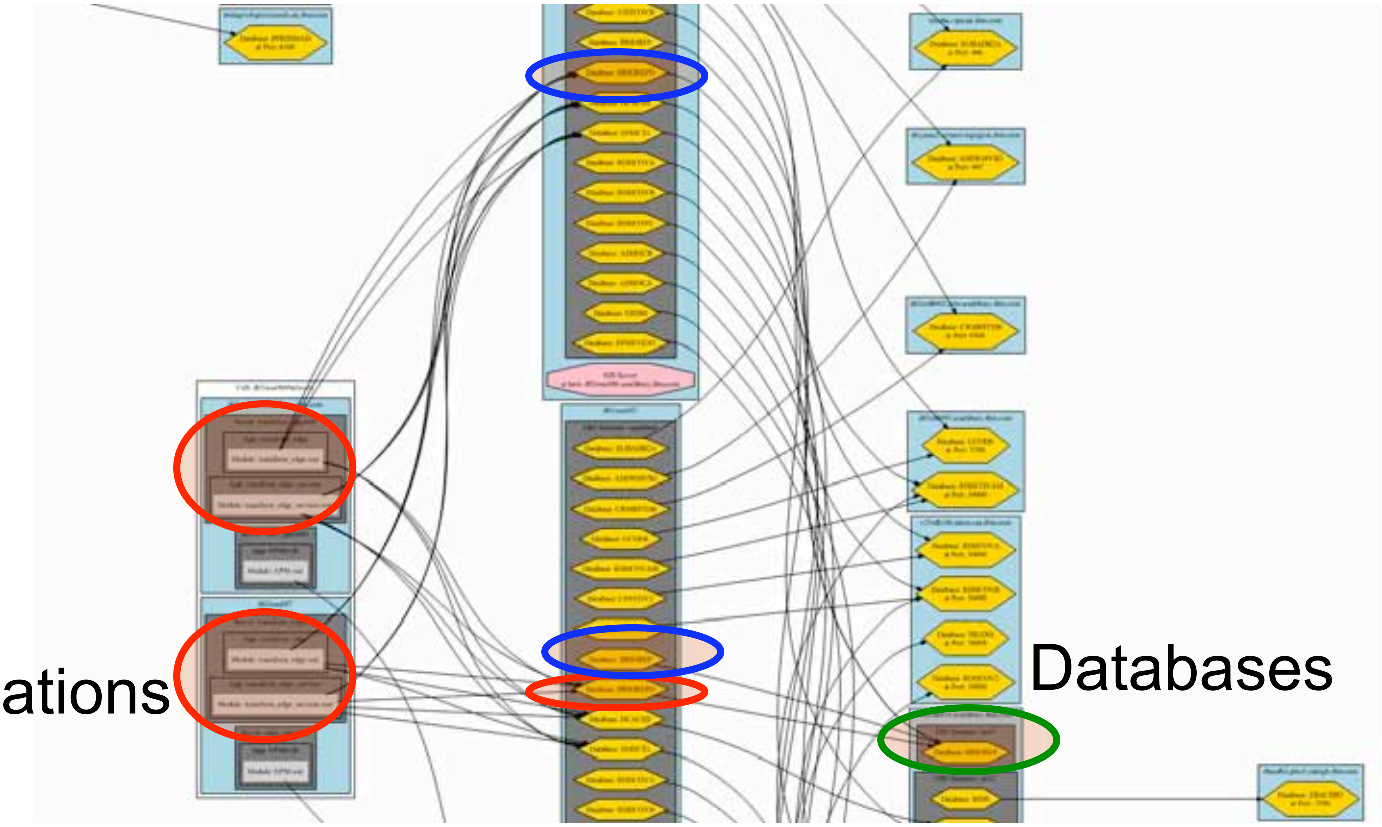

Applications

Databases

Databases

Conclusion

- **Automated discovery technologies are here to stay**
- **Galapagos unique for its ability to discover deep app-data and data-data dependencies**
- **Particular relevance to availability management functions**
  - Recovery Planning
  - Identifying Actually Enforced Failure Boundaries
- **Future work: Automate and integrate availability management functions with Galapagos**