

# Introduction to the DICE Program Through Case Studies

*"Even a mistake may turn out  
to be the one thing necessary  
to a worthwhile achievement."*

*Henry Ford*

**Al Stutz, CIO**

**Rick Beam, Senior Project Manager**



# About the DICE Program

- [Avetec.org](http://Avetec.org)
- Not-for-profit public benefit organization
- Founded and located in Springfield Ohio in 2005





# Nationwide Test Environment for Technology Innovations



# Researching, Testing & Accelerating Innovations

With

*internal testing,*

how do you mirror the real production environment of diverse, geographically distributed data centers?

*The Reality...*

**Diverse technologies are being used in data centers that are geographically dispersed across dissimilar architectures making realistic testing even more critical**



A report from the Systems Sciences Institute at IBM states that the cost to fix an error found *after* product release was four to five times as much as one uncovered during design, and up to 100 times more than **“What is the true cost of correcting problems after deployment?”** problems after deployment phase.

***“Have you ever installed a product and it did not work?”***

***“Have you ever lost money or credibility on a software or hardware product because it was not fully tested?”***

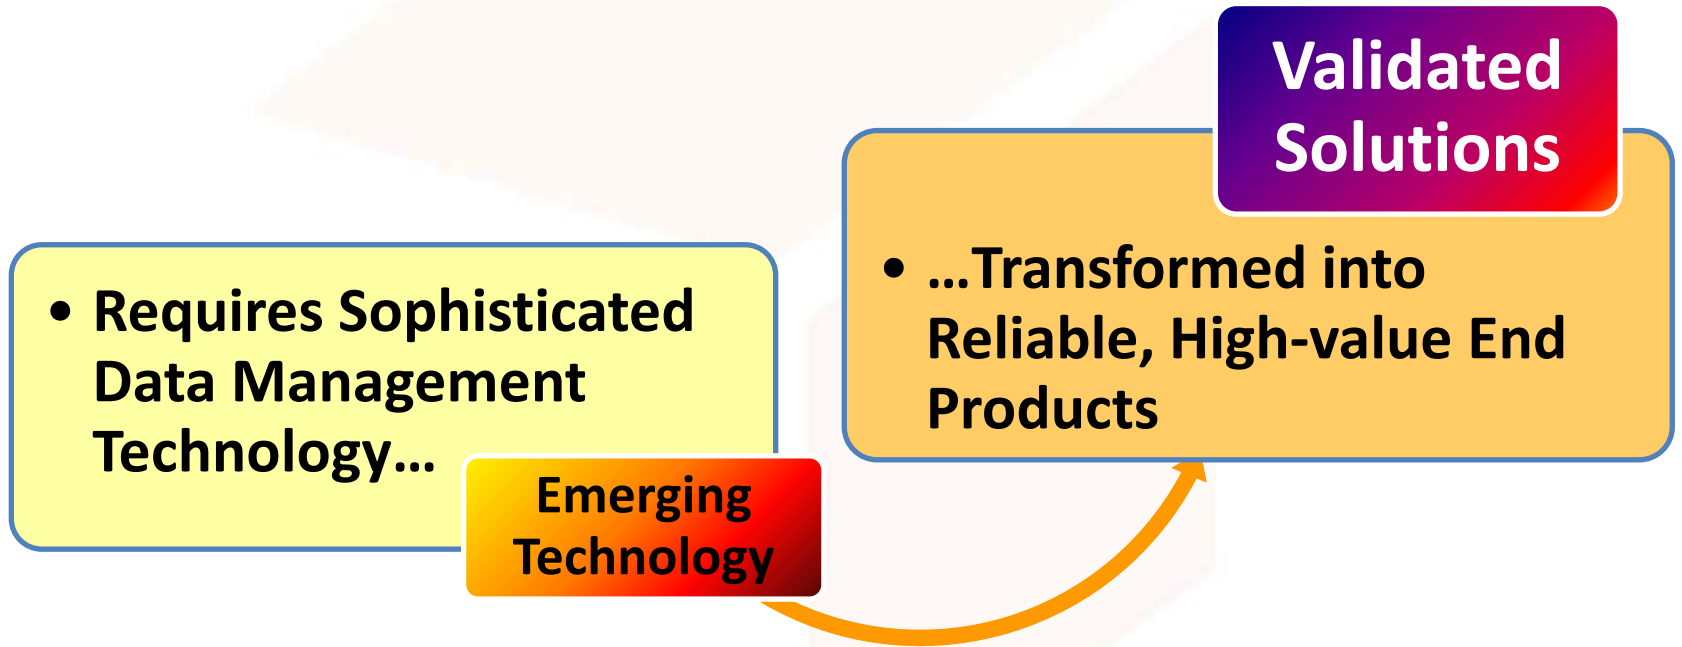
***“Have you ever been forced to delay deployment of a solution due to defects?”***

***“What is the true cost of correcting problems after deployment?”***

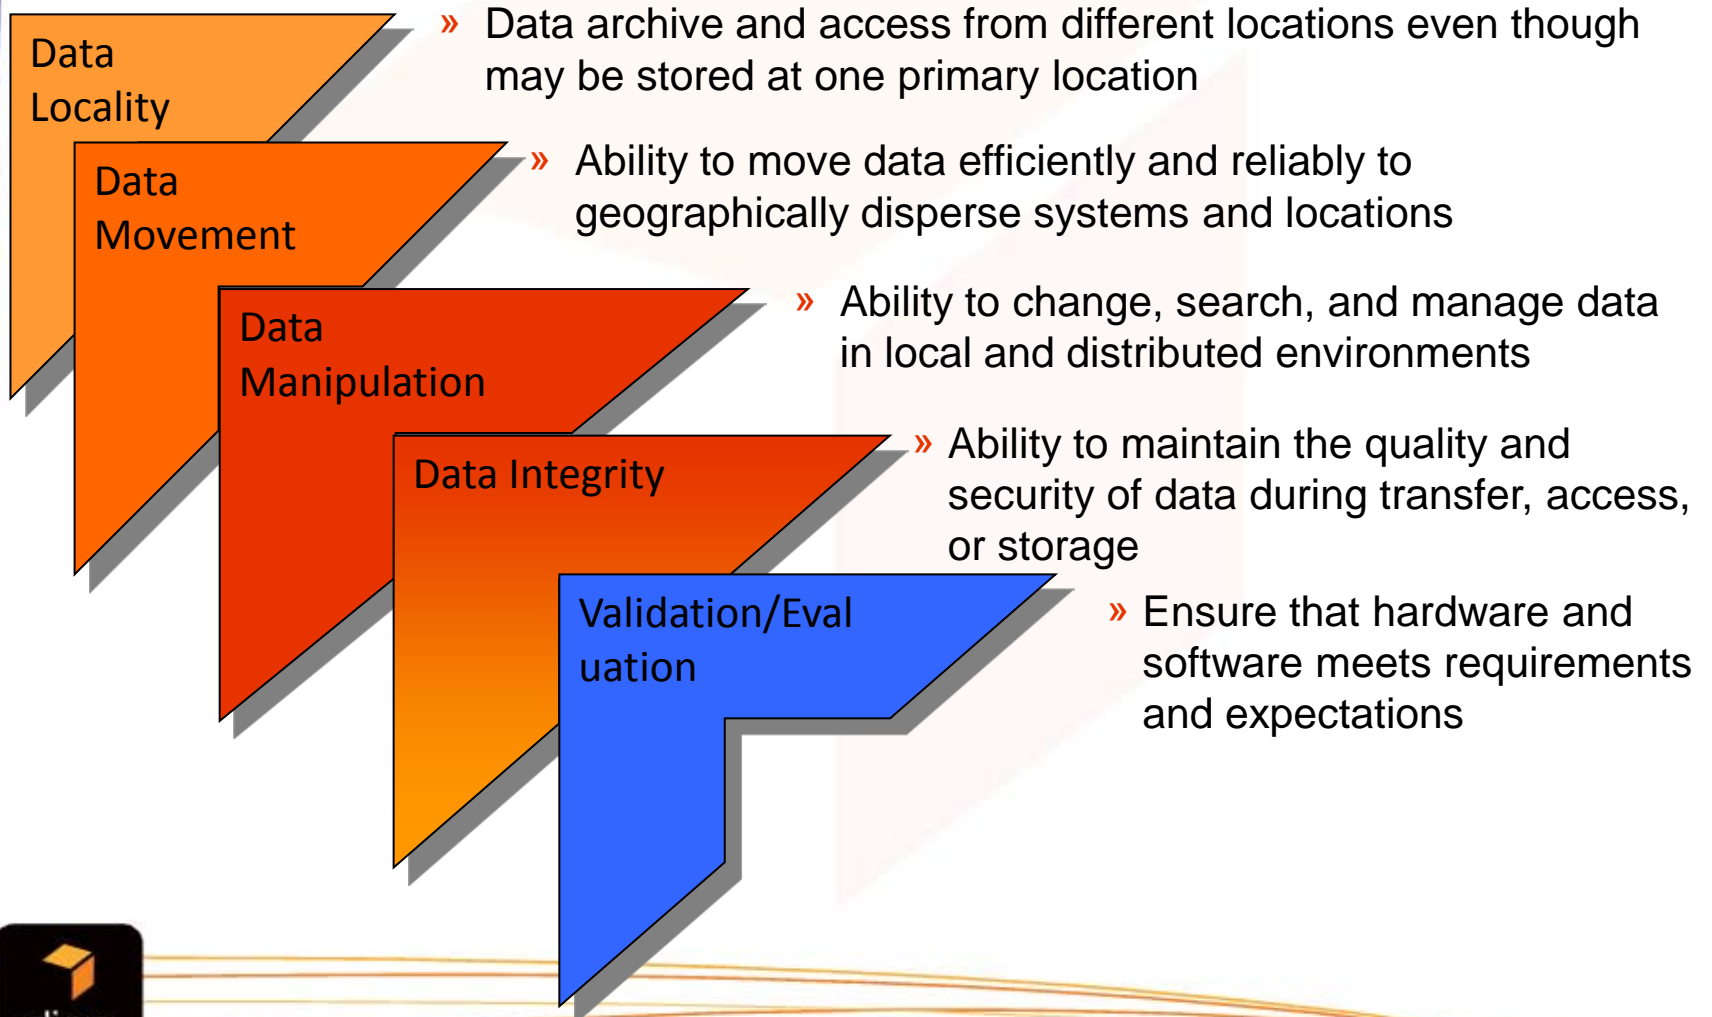
# ...And I'm Tired of the Answers

- *Who in their right mind would have that configuration?*
- *What in the world were they thinking?*
- *We never expected anyone to do that.*
- *We need to fix the documentation.*
- *This is an undocumented feature!*
- *If they don't like our product, let them use our competitor's.*

# The Environment is Constantly Changing




# DICE Challenge Areas



# Researching, Testing & Accelerating Innovations

**DICE, the  
community's  
independent  
resource  
for...**



# The Value: *For Solution Providers*

Plan

- ◆ Focus R&D investments on critical data management challenges
- ◆ Understand requirements of HPC user community

Design

- ◆ Optimize product design, performance & functionality
- ◆ Correct deficiencies & vulnerability issues before deployment

Test

- ◆ Test technology based on unbiased, independent evaluations
- ◆ Receive technology validation certificate

Market

- ◆ Enhance validity & marketability of emerging solutions
- ◆ Accelerate time-to-market

**Reduce Risk**

**Improve Bottom Line**



# The Value: *For Centers*

## Plan

- ◆ Participate in establishing benchmarks and metrics
- ◆ Identify critical data management challenges

## Program

- ◆ Invest in validated technologies
- ◆ Validate product concept and reliability

## Test

- ◆ Accelerate deployment of new systems, tools & enhancements without disrupting production environment
- ◆ Implement validated solutions based on a realistic test environment

## Discover

- ◆ Expedite scientific breakthroughs from data analysis to discovery
- ◆ Expand HPC knowledge and focus resources on core mission
- ◆ Receive unbiased evaluation reports

**Reduce Risk**

**Improve Cost-  
Effectiveness & Service**



# Our Alliance

- Adaptive Computing
- DoE
- DoD
- AMD
- BlueArc
- Cisco Systems
- CSC
- DataDirect Networks
- Dell Computer
- EMC<sup>2</sup>
- FalconStor
- Force 10 Networks
- Foundry Networks
- General Atomics
- Hitachi
- IBM
- IBRIX
- Intel
- IRODS
- Isilon System
- Juniper Networks
- Brocade/McData
- Mercury
- Microsoft
- NASA
- NetApp Storage
- Novell
- Obsidian Research
- Ohio State University
- Ohio Supercomputer Center
- Oracle
- Panasas
- Paraview
- Pillar Data Systems
- Platform Computing
- QLogic
- Red Hat
- SGI
- Sun Microsystems
- Teradata
- University of Maryland
- ViON
- Wolfram Research



# DICE Technical Advisory Panel

- Comprised of HPC leaders
- Provides external & end-user perspective
- Identifies critical data management & other challenges

**the TAP**  
technical advisory panel

Dan Duffy, NASA 

Chris Jordan, TACC 

Steve Conway, IDC 

Matt Leininger, LLNL 

Lee Ward, Sandia 

Kevin Wilson, P&G 

Kevin Regimbal, PNNL 

Ron Bewtra, NOAA 

Lloyd Slonaker, AFRL 



# Community-Driven

- **Technical Advisory Panel**
- **Government agencies, Vendors, companies**
- **Open process**
- **Publically available reports**

***The DICE program is our community's UL service***





**Independent, Unbiased 3<sup>rd</sup> Party  
Evaluations**



# Case Study #1

## Study of Power and Cooling at Data Centers

**Joint Study by IDC and DICE Program**



# The Challenge


- Data centers, including High Performance Computing (HPC), enterprise & telecommunication, are critical resources for competing in a world economy
- A 2007 EPA Report to Congress indicated that unless energy efficiency improvement trends are not accelerated, energy consumption by data centers could nearly double again in five years
  - By 2011, the energy consumption could reach 100 billion kilowatt hours (kWhs)
  - This rate of increase could require building two additional power plants per year in order to provide the required power *perhaps MORE*
- The demand for computer resources has led to a significant growth in the number of data centers

# The Challenge (Cont'd.)

- In some areas of USA and Europe, data centers can not obtain additional power
- The explosive growth in the number and size of data centers is creating an unsustainable demand for power requirements
- The increased power requirements have led to a situation in which new data centers are being built primarily to increase power capacity or exploit lower local power costs, not because the current centers are running out of floor space


# Avetec & IDC 2-part Research Study

- Study #1: *Power & Cooling Practices and Planning at HPC Data Centers* -- Spring 2009
- Study #2: *Power & Cooling Solutions for Enterprise and HPC Data Centers* -- October 2009



REPORTS ARE AVAILABLE TO  
ANYONE WHO IS INTERESTED!

[WWW.DICEPROGRAM.ORG](http://WWW.DICEPROGRAM.ORG)




# Research Study: *Objective*

- Examine power & cooling issues, practices & plans of data centers on a worldwide basis and across all major categories of centers (HPC, enterprise, and communications)
- Explore the critical infrastructure issues facing data centers
- Identify and assess current practices and approaches, as well as plans to improve data center energy efficiency
  - Especially cooling for computer systems in use today and planned over the next 2 to 5 years

# Survey Basics

- Survey includes 49 respondents representing 223 data centers as well as vendors
  - Intent was to focus on data centers from industry, academia and government

**Data Center Size**



- Respondents were from US, Europe and Asia

# Key Findings: *Internal Organizational Practices*

## ➤ ***Who Controls & Pays the Power Bill?***

- Typically, not the CIO – less than half the centers PAY for costs directly
- Less than half track and report power & cooling costs
- The HPC & IT community needs more insight and accountability for the total cost of ownership (TCO) to include power and cooling usage

## ➤ ***“If You Can’t Measure It, You Can’t Manage It”\****

- Only 34% of the centers use specific metrics and/or statistics to measure and report energy efficiency
- Few centers seriously track power and cooling efficiency
- However, 49% plan to improve the measuring and tracking of energy efficiency in their centers over the next 18 months

(\*Deming & Drucker)

# How Data Centers Currently Track & Measure Power & Cooling

- **Q: Describe how you track and report the energy efficiency of your data center? (selected responses):**
  - *Not to good at it.*
  - *Only track with usage, utilization of the power supplied to the data center.*
  - *We monitor invoices and consumption rates.*
  - *We track the power usage through the PDUs and the power supplied to the CRAC units. At this point energy efficiency is not reported.*
  - *We calculate our costs by multiplying current energy rates with specs provided by respective vendors. It's a very manual process.*
  - *Return on investment, IT performance, and facility utilization metrics.*
  - *We monitor our electric bill, A/C cycling and temperature, and UPS load.*

***Are You Ready for a CAP & Trade Law?***



# Key Findings:

## *Internal/Organizational Practices*

### ➤ *Mandates to Improve:*

- Majority of sites (71%) operate under relatively weak internal or external *mandates* to reduce energy consumption, but it is a top concern to data centers
- Only 29% of the data centers have a formal roadmap for improving power and cooling
- Whereas, 75% vendors have formal roadmaps for helping data centers achieve greater energy efficiency
- The Largest Non-Technical Barrier to Increased Energy Efficiency Is *Financial*
- Data centers see potential benefits, but also drawbacks in industry-standards

# Key Findings: *Data Center Consolidation & Building*

- Power & cooling infrastructure limitations were the biggest barriers to increasing available HPC resources
- Nearly all data centers (96%) consider “green” design criteria important to their HPC system and data center planning process
- Half of the data centers are considering distributing their resources to multiple buildings or locations
- Approximately 67% of the sites are planning to expand or build new HPC data centers
- Data centers will increasingly be located in areas with the lowest power costs

# Key Findings: *Need for Evolutionary, But Verified Solutions*

## ➤ *Different Perspectives:*

- 62% of vendors expected game-changing power and cooling technologies in the next five years
- Only 36% of user sites have these expectations

## ➤ *DOE Roadmap Recommends Independent, Neutral Testing Body*

- Most Data Center Managers Rely On Vendor-Supplied Specifications
- Most Data Centers Favor the Idea of An Independent Test and Validation Service
- 88% of vendors Are "Very Likely" Or "Somewhat Likely" To Use An Independent Test And Evaluation Body If A Suitable One Were Available

# Case Study #2

## Information Life Cycle Management

**Study completed for  
the Defense Modernization Program**



# The Challenge

**The Modernization Program was interested in considering ILM/DLM products to manage end user data**

**The DICE program was asked to assist with:**

1. Define the requirements/features the DoD program should consider
2. Test four vendor products with regard to:
  - Ease of Installation
  - Usability
  - Feature Set
  - Functionality in a Distributed HPC Environment

# Formal Requirement Gathering

- Requirements were gathered based on the customer's need or problem.
- Best practices across the community were examined
- Vendors were approached to understand current and emerging technology
- The methodology for the evaluation through a weighting process each of the requirements in terms of function, operational impact, and other requirements
- The customer signed a formal document

# Testing the Products

- Each product chosen to be evaluated is evaluated on the same requirements.
- Test case scenarios are established based on the approach of the product and how it functions to complete the requirement.
- Testing includes every aspect on implementation.
  - Installation and Configuration
  - Maintenance
  - Troubleshooting
- Any discovered undocumented information or caveats are reported

# Reporting

- Each product will have its own separate UNBIASED report
- No Product comparison
- Only reported in relationship to the requirements
- Product comparison is only done by the customer by comparing reports
- The Vendor also benefited from an evaluation from not only fixing

# How were the results of the study used?

1. Developed a Request for proposal based on the requirements
2. Used the study results to challenge the vendor community to solve many of the issues

DoD selected Nirvana from General Atomics. They are in the implementation phase

## NIRVANA Storage Resource Broker (SRB)

- Tested and Validated Product Functionality
- Collaborated on Product Enhancement Ideas
- Verified and Enhanced Installation Procedures
- Collaborated on Enhancing Documentation

**A Win-Win for the Community**

*“The DICE test bed evaluation was the number one step in gaining the customer’s confidence in our solution”*

- Greg Granello, Nirvana Market Development

# Case Study #3

## Gauging Parallel File System Performance: Benchmarking and Evaluation

*Current Project*

# The Challenge

- **DOE Sandia funded the DICE Program to initiate phase one of a project to identify community needs and interest in an open source suite of benchmarks for parallel file systems evaluation that would lead to a fair head-to-head comparison**
- **Assist data centers in selecting the best parallel file system for their job mix**

# Project Team

- **Tracey Wilson, CSC, Project Lead**
- **Lee Ward, Sandia National Lab**
- **Ed Wahl, Avetec**
- **Paul Buerger, Avetec**
- **Armen Ezekilian, Avetec**
  
- **DICE TAP Involvement:**
  - ◆ Dan Duffy, NASA Goddard
  - ◆ Matt Leininger, Lawrence Livermore National Lab

# Parallel File System *Issues!*

- **Lack of standardized metrics for performance**
  - Need exists for file system performance comparison
  - Need exists for defined benchmarks for local and remote file systems
  
- **Scaling:**
  - Scaling of current file systems has unpredictable performance impacts
  
- **Load:**
  - Performance varies in file systems with loads of 50-70%
  - Fragmented data is a problem

# Project Goals

## ➤ **Parameterized Benchmark Suite**

- Simulate many different workloads
- Expose more tuning/configuration options in the benchmarks
- Suite Extensible so users can add new benchmarks
- Have unique benchmarks for synthetic workloads
  - Transactions
  - Streaming
  - Random I/O
  - Various Read/Write Ratios


## ➤ **Normalization is needed between differing architectures**

- Need a formula
- Aging and fragmentation of data should be considered


# First Task: Survey the Community

- **What are the file systems in use today**
- **What percentage of the storage in HPC Research centers is devoted to parallel file systems**
- **What are the current measuring tools and what are the limitations**


# Percentage of Spinning Disks Used for Parallel File Systems



# Data Center Use of File Systems



# Benchmarking Tools Used



# Need for Normalization

- **Normalization of different file system architectures**
  - Over 60% say absolutely yes
  - About 25% are uncertain
  - Only a small portion disagree with the need

# Step 1: Data Collection

## ---- I/O Characterization

- **Collect Trace and Benchmark results from sites across the globe. Maintain site confidentiality**
- **Parse the data to identify patterns**

# Step 2: Benchmark Framework

- **Design framework tool to evaluate file systems**
  - Parameterization is a key
  - Summary of current tools and their capabilities
  - Must be extensible in order to add new file system architectures

# Parts 3, 4 & 5

3. Present results to the community and seek advice on the Benchmark Framework
4. Develop open source suite of codes
5. Develop normalization formula

# How *YOU* Can Participate!

- **Contribute *YOUR* Trace Data**
  - **Benchmark and Trace data**
  - **Collection Site <http://webproj.usa40.net>**

# Where to Find Our....

Research, Case Studies, Reports and  
Whitepapers

[www.diceprogram.org](http://www.diceprogram.org)

- Information Lifecycle Management (ILM)
- Data De-duplication
- Power and Cooling of Large Data Centers
- Evaluation Reports & Studies



# Quite Simply...

**Increase Success**

**Reduce Risk**



# For More Information

Rick Beam, Senior Project Manager

[rbeam@avetec.org](mailto:rbeam@avetec.org)

937-322-5000 x2024

Al Stutz, CIO

[al@avetec.org](mailto:al@avetec.org)

937-322-5000 x2025

