OSIsoft PI High Availability (HA) & Virtualization

Complementary & contrasting technical approaches to HA

Presented By:

Guy F. Acciai
Integrated Manufacturing System Architect
Corning Incorporated - Blacksburg, VA

OSIsoft®

Agenda

- Who is Corning and what do we do…
- Corning Life Science (CLS) manufacturing
  - Business need and value of High Availability (HA)
- HA & Virtualization
  - PI and VMware product offerings – Pros/Cons & Tradeoffs
  - CLS - Final Architecture
  - Usage of complementary tools & “watch-outs”
- Usage to date
- Futures Needs & Questions
  - PI futures
  - VMWare futures

CORNING

OSIsoft.

Agenda

- Who is Corning and what do we do…
- Corning Life Science (CLS) manufacturing
  - Business need and value of High Availability (HA)
- HA & Virtualization
  - PI and VMware product offerings – Pros/Cons & Tradeoffs
  - CLS - Final Architecture
  - Usage of complementary tools & “watch-outs”
- Usage to date
- Futures Needs & Questions
  - PI futures
  - VMWare futures

Corning Incorporated

Founded: 1851

Headquarters: Corning, NY

Employees: ≈23,500 (worldwide)

2009 Revenues: $5.4 Billion

Fortune 500 Rank: 391

- Corning is the world leader in specialty glass and ceramics.
- We create and make keystone components that enable high-technology systems for consumer electronics, mobile emissions control, telecommunications and life sciences.
- We succeed through sustained investment in R&D, over 150 years of materials science and process engineering knowledge, and a distinctive collaborative culture.

OSIsoft®

Corning's Culture of Innovation

1879 Glass envelope for Thomas Edison's light bulb

1915 Heat-resistant Pyrex® glass

1934 Dow Corning silicones

1947 Processes for mass producing the television bulb

1952 Glass ceramics

1960 Fusion draw process

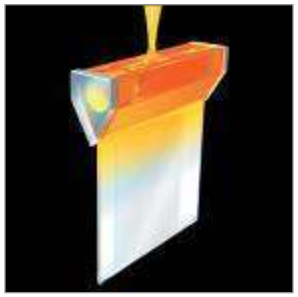

1970 Ceramic substrates for automotive catalytic converters

1972 First low-loss optical fiber

1984 High-throughput label-free screening for drug discovery

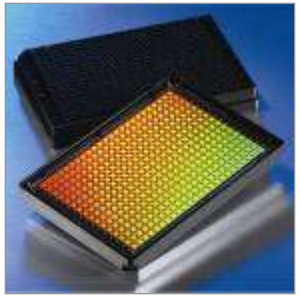

2006 AMLCD glass for TVs & computers

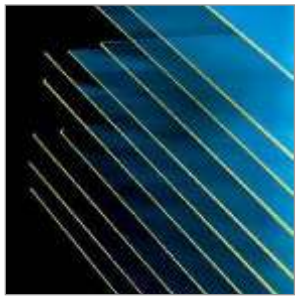

OSIsoft®

Market Segments & Operations

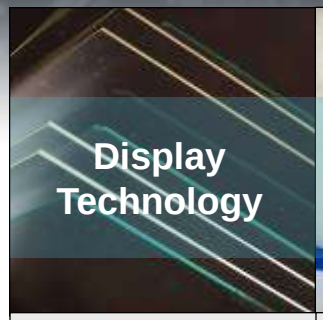

Display Technology

- LCD Glass Substrates
- Glass Substrates for OLED and LTPS-LCD

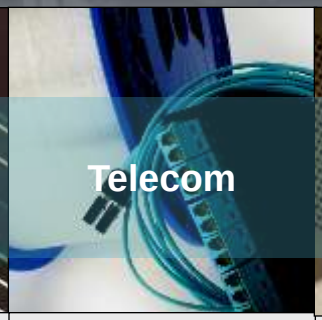

Telecom

- Optical Fiber & Cable
- Hardware & Equipment
  - Fiber optic connectivity products
  - Optical connectivity products

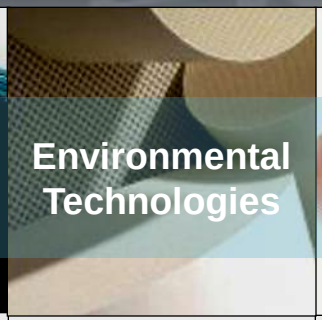

Environmental Technologies

- Emissions Control Products
  - Light-duty gasoline vehicles
  - Light-duty and heavy-duty diesel vehicles
  - Stationary

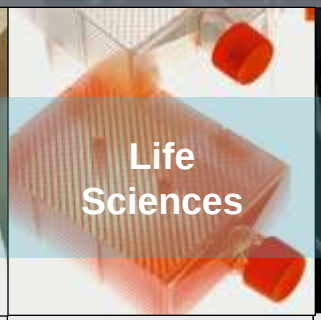

Life Sciences

- **Cell Culture & Bioprocess**
- **Assay & High-Throughput Screening**
- **Genomics & Proteomics**
- **General Laboratory Products**

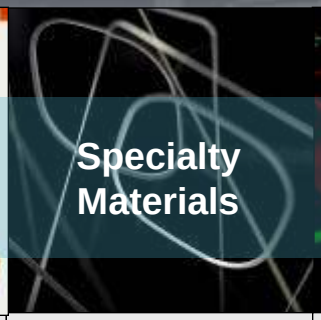

Specialty Materials

- Display Optics & Components
- Semiconductor Optics/Components
- Aerospace/Defense
- Astronomy
- Optical Metrology
- Ophthalmic
- Telecom Components
- Specialty Glass

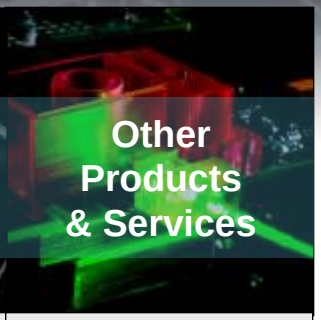

Other Products & Services

- Display Futures
- New Business Development
- Drug Discovery Technology
- Equity Companies
  - Cormetech, Inc.
  - Dow Corning Corp.
  - Eurokera, S.N.C.
  - Samsung Corning

Corning Life Sciences (CLS)

Developer, manufacturer and global supplier of innovative life science product solutions for over 90 years

OSIsoft®

CLS Manufacturing

- Over time, products have shifted from borosilicate glass to greater use of sterile/disposable plastic
- Extensive use of injection molding to form plastic parts used for medical testing and drug discovery
- FDA regulation and traceability requirements
- Focus on overall equipment & labor effectiveness (OEE) as well as scrap & downtime pareto analysis
- Business decision to deploy a Manufacturing Execution System (MES)

Business need for HA

- Use of existing PI server as SCADA (and historian) to act as IO interface to MES
- MES generates product labeling & interacts with inventory systems. Availability is paramount

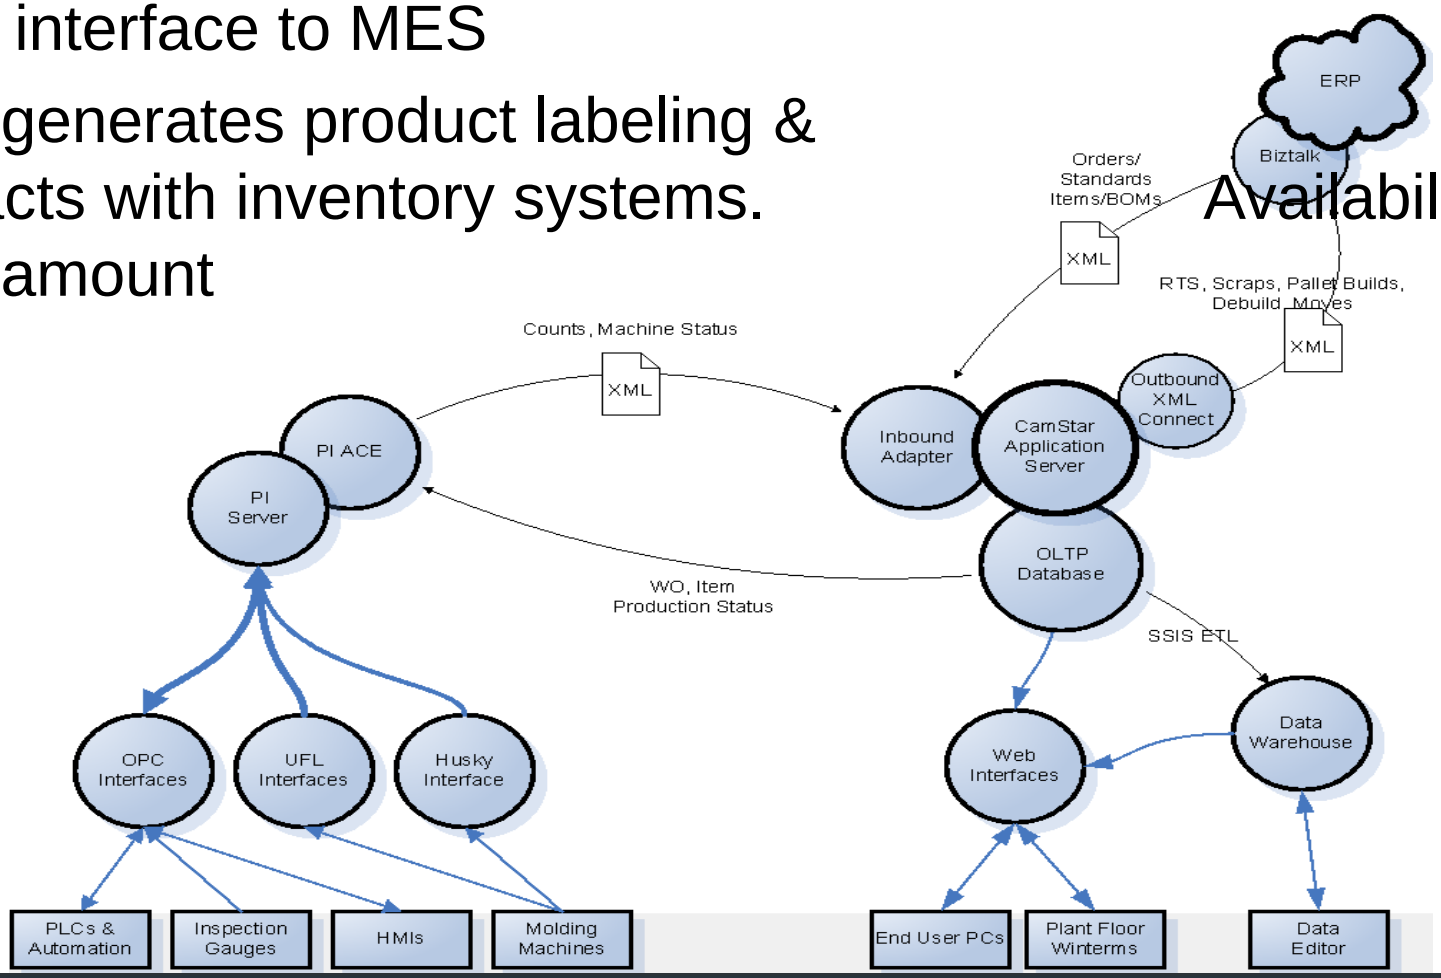

High Availability

- Definition: *a system design & associated implementation that ensures a certain degree of operational continuity during a given time period*
  - Design for availability (given planned downtime & unplanned failures)
- Major approach is replication & redundancy
- Historical Approaches:
  - Clusters & load balancing
  - Redundant/Failover Hardware (ie, Tandem, Status, RAID)
  - Marathon EverRun/Stratus ftServers (with ActiveUpgrade/Avance SW)
- OSIsoft History
  - OSIsoft PI server on a MSCS
  - Interface Redundancy & PI Server HA

PI Interface Failover

- Been available in some interfaces for several years
- OPC is “model interface”
  - Need to “license two OPC” interfaces, or rule of 3 interfaces
- Supports two “types” of HA
  - Support for “redundant” OPC servers
  - Support for “redundant” PI OPC interfaces
    - UnInit based (preferred) or MS cluster based
- Our implementation, uses redundant OPC servers on redundant interface nodes

PI Interface Failover Architecture

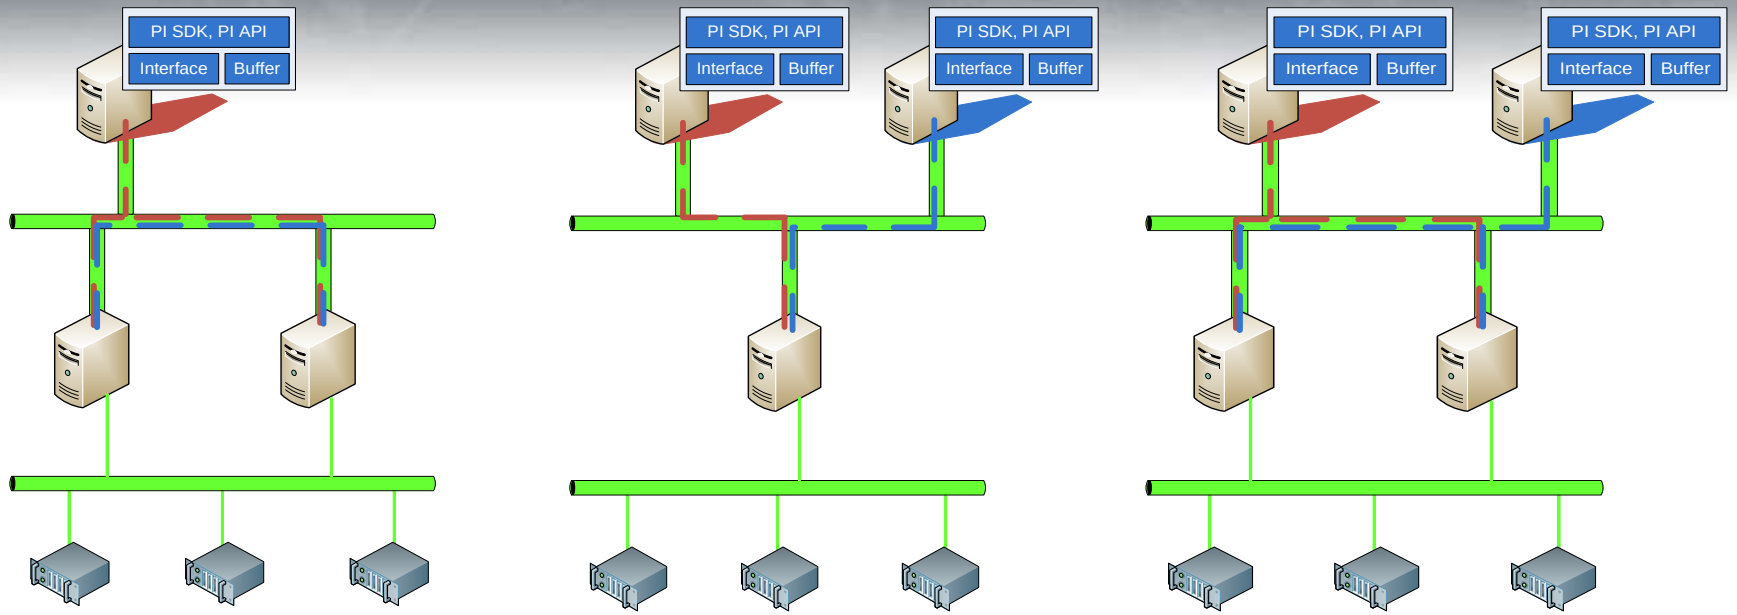

- OPC server coordination support for 3 different modes:
  - Hot Mode, Warm, Cold
- Redundant OPC Interface Coordination
  - Use of “interface control” tags - to avoid the “Split-Brain” problem

OSIsoft®

Redundant Interface Coordination

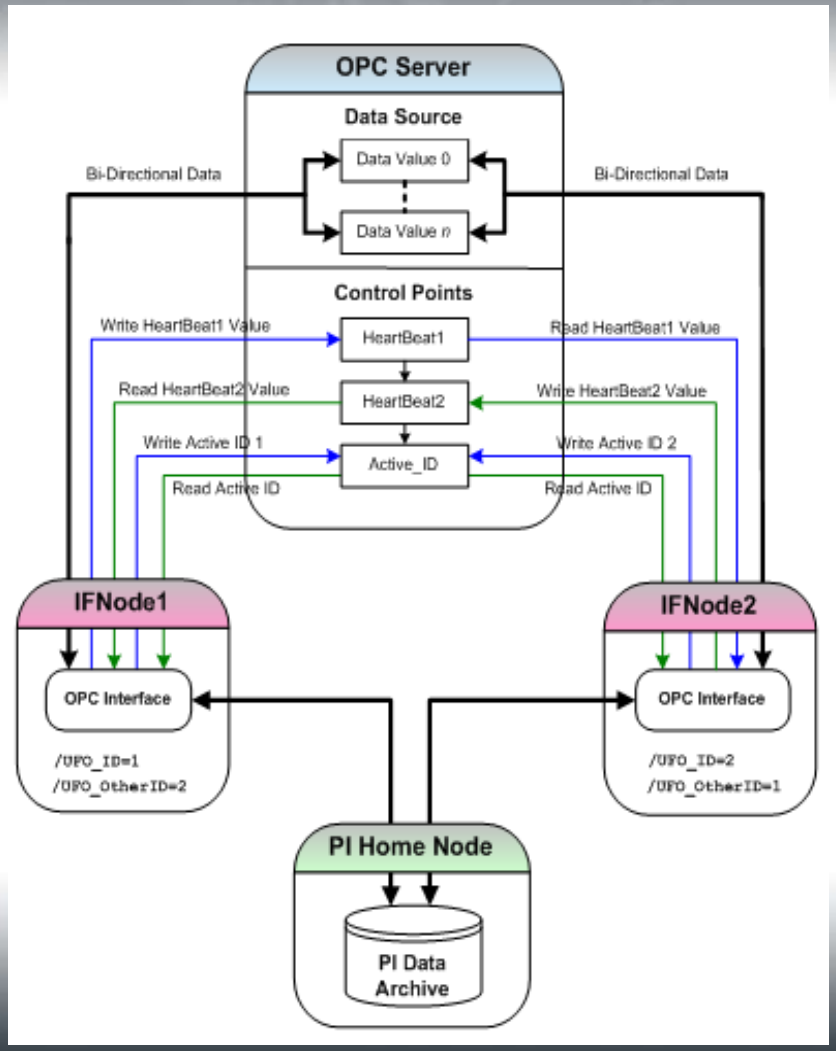

- Avoid the “Split-Brain”
  - Control which interface is talking to PI or the OPC server
- Done thru PI server with 8 specific “control” tags
- Alternate communications path via:
  - Phase I (thru OPC server)
  - Phase II (thru a shared file, just on a file server) – Our approach

OSIsoft®

PI Server Redundancy
…The *Collective*…

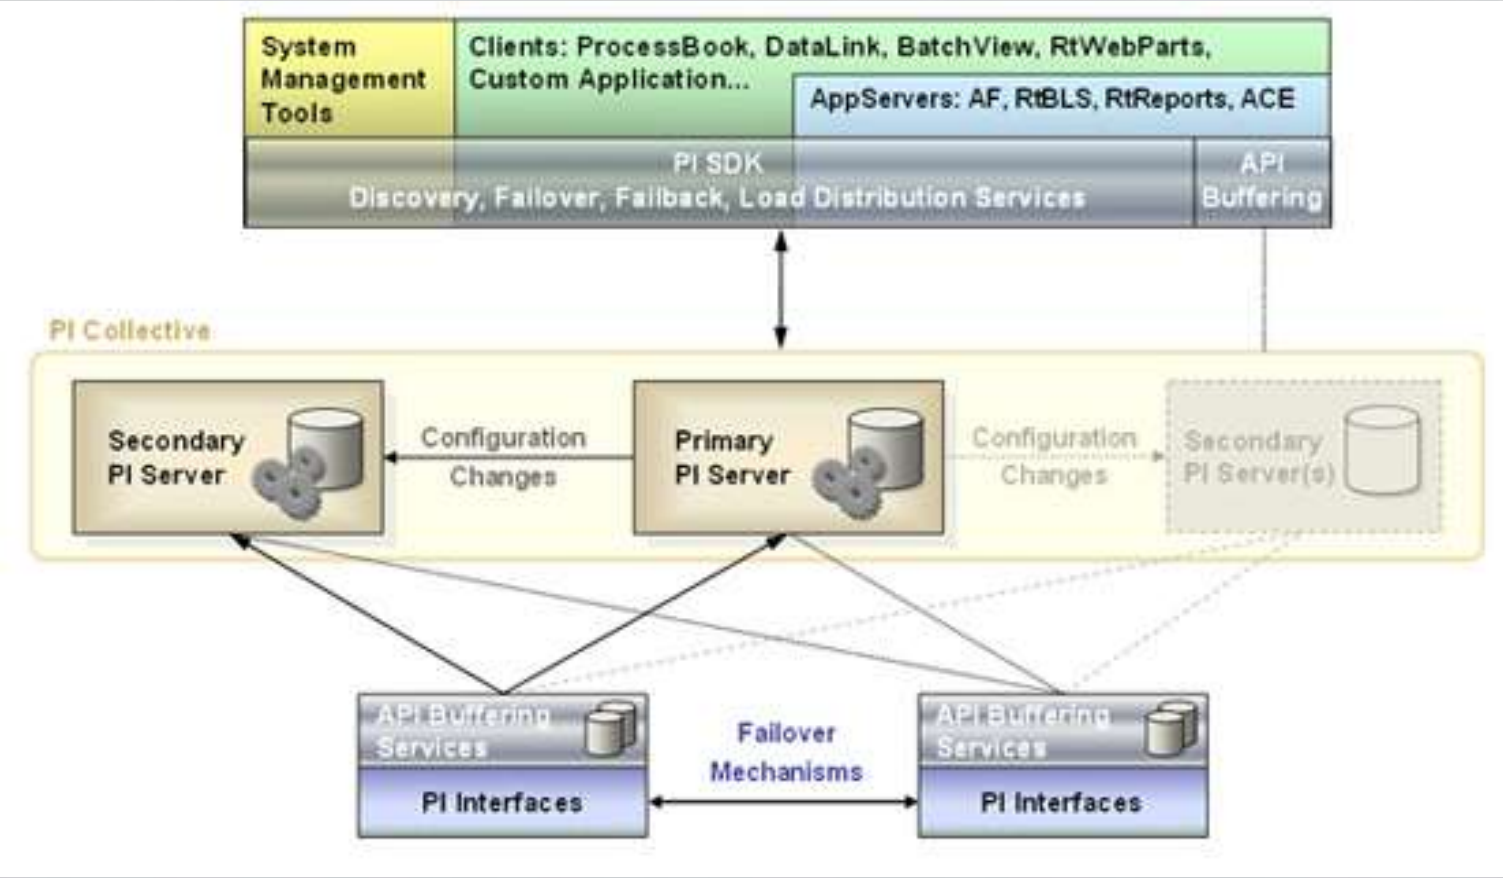

OSIsoft®

CLS HA Topology

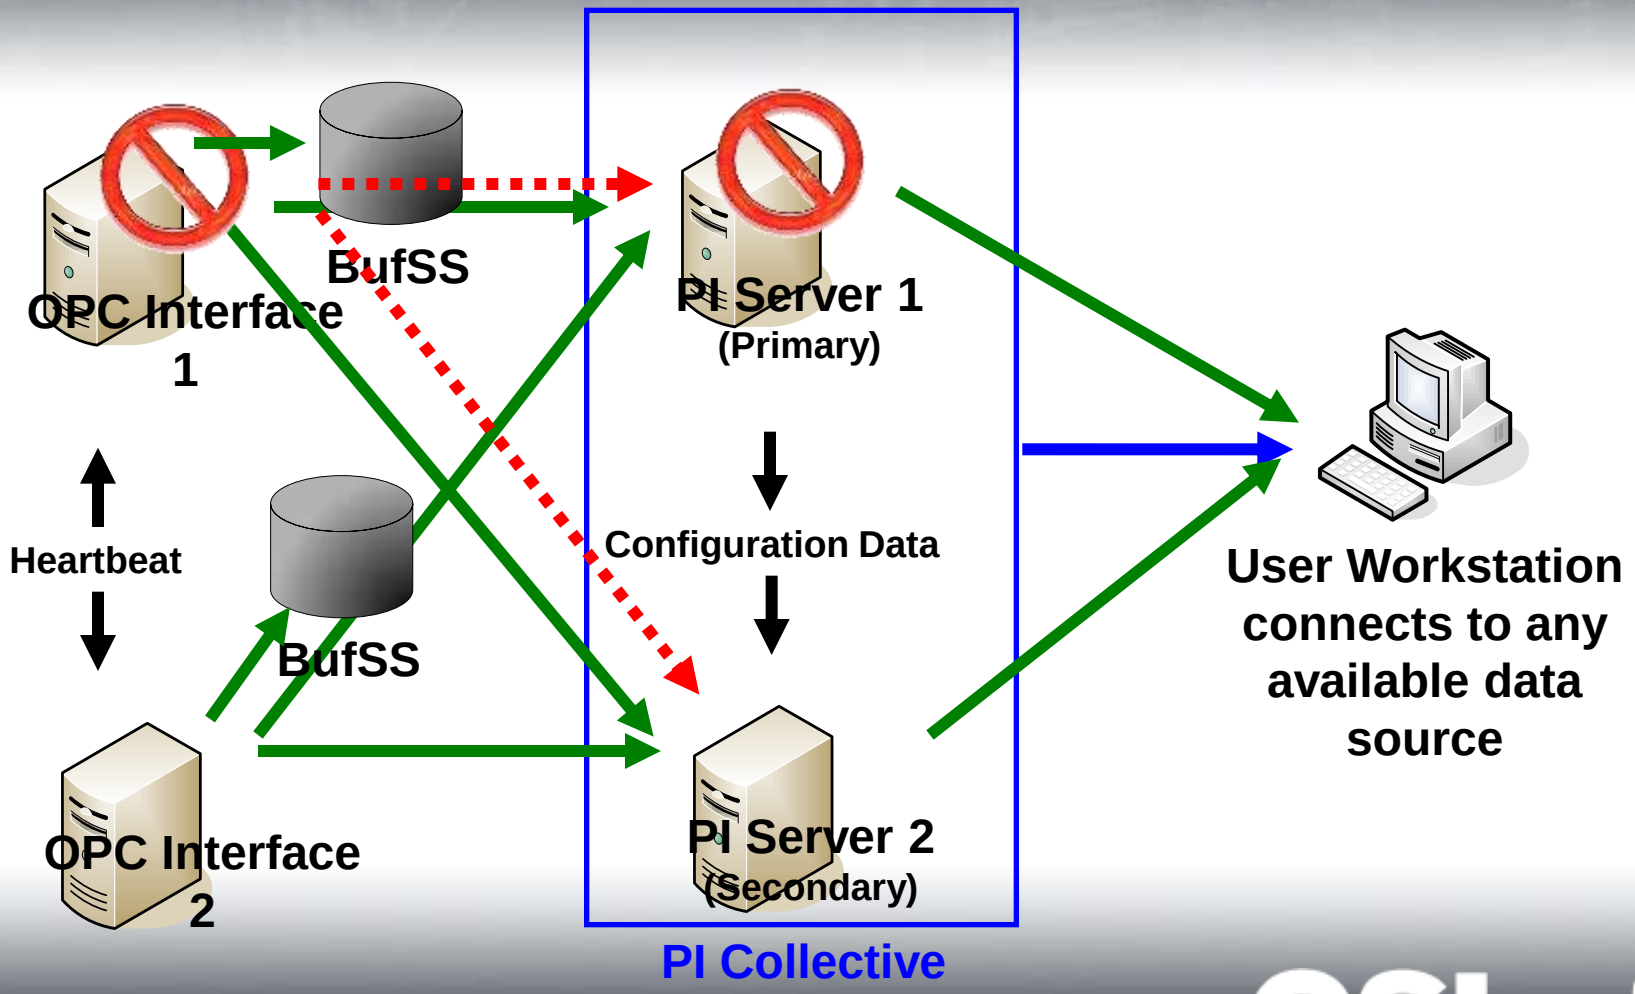

Collective Naming & Migration

- For “migration” of a new “collective” to replace an existing PI server – Must name collective with same name as old server
  - Collective is virtual - Has no IP address & only “visible” to PI SDK/API
  - Need to match/force “PI Server ID”
    - Some SDK/API’s tools use ServerID, and others Server Name
    - Tools are good, and migration is not hard
      - But fairly complex topic, and requires understanding and planning
  - Migration also involved network team to control “timing” around DNS switch to point old server name to new collective
  - Our Upgrade Approach
    - Interfaces first, to support buffer to “old server” & collective for testing
    - Collective Formation (Keep old server for fallback)
    - Client Upgrade
    - Finally, shutdown old server

OSIsoft®

PI HA - Pros and Cons

- Pros - Improved:
  - Reliability (semi-transparent failover)
  - DR support (multiple servers in diverse locations)
  - ***Application availability during maintenance (rolling upgrades & security patching)***
- Cons - Higher Costs and Complexity
  - SW OSIsoft - (Additional PI Server and interfaces)
    - Enterprise License can help mitigate this
  - HW – Duplicate HW (Servers AND Storage)
    - Virtualization helps mitigate some of these costs
  - Greater complexity (Greenfield easy, Migration more complex)
  - Some interfaces lack UniInt failover & limitations around HA V1

Current PI HA Limitations:

- **Non-Interface data (Manual writes/Lab Points) are not currently propagated to the collective**
- PIBagen will not generate the same batch-ids for batches generated on different servers. PIBagen and “old” batch data is not replicated between servers in the collective
- Changes to tags/metadata can ONLY happen on primary
- Archive annotations and changes to annotations are not currently propagated to the collective.
- Acknowledgement of alarms is not currently replicated between servers in the collective.
- Many of these are to be addressed in PR2 of HA

Virtualization

- Definition: *Virtualization is abstraction performed on a given hardware platform by host software (a control program), which creates a simulated computer environment, a "virtual machine", for guest software.*

- Alternatives – In the Server Space (on Windows)
  - MS Hyper-V (and MS Virtual Server)
  - VMWare ESX Cluster + VMotion (and VM Server)

- Problems its solves
  - Server consolidation - reduced energy and physical footprint
  - Improved availability & operational flexibility (provisioning)
  - Disaster recovery

Corning's drive to virtualization

- Large multi-year effort by central IT groups to setup failover ESX clusters in each site to support "standard" IT functions (4 servers):
  - IT Core Services (File & Print/Backup/Patching/AV/DC/DNS/DHCP/etc.)
  - VMWare enterprise license
- General approach is to "virtualize everything we can"
- Attempt to "leverage" this standardized investment for manufacturing and site specific applications
  - PI is a "well behaved" application, with predictable resource demands, that is easily virtualized

OSIsoft formally supports VMs

- ArticleID: 3062OSI8 - **A PI Server is supported when installed on a virtual machine that's running a supported operating system**.
  - **With the caution**: *Considerations must be made to evaluate how much resource contention is expected at peak loads. Critical resources are usually I/O bandwidth and disk resources*

- Corning has multiple PI servers running on VMware – Fairly painless to virtualize, given sufficient resources (generally scales linearly with tag count & polling rate)

VMWare Architecture

- ESX physical clusters
  - HA – Auto-Restart (on OS or HW failure – Not Application)
  - DRS – Load Balancing
  - Implies SHARED Storage
- Key component to Corning's HA environment is VMotion
  - Move ***running*** virtual machines from one physical server to another ***with NO impact*** to end users

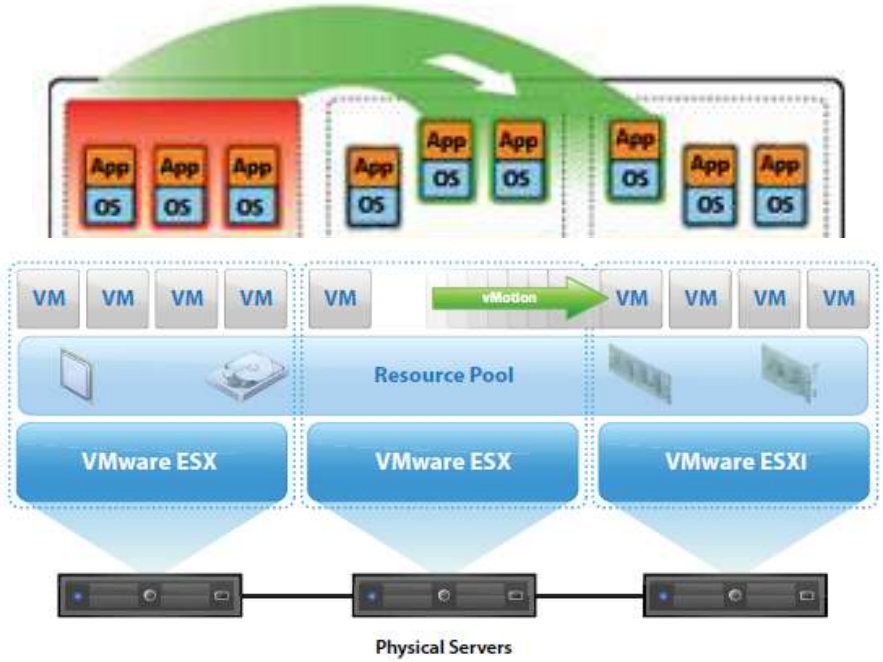

VMWare - Pros and Cons

- Pros - Improved:
  - Reliability (automatic restart on “crash”)
  - Scalability & QOS (Easy “add” of additional resources to guests & “hot” additional of nodes to ESX cluster, DRS load balancing)
  - DR support (VMware Snapshots & Distributed Clusters)
  - Availability during planned HW maintenance (Vmotion)
  - Better utilization of HW, energy & floor space savings
- Cons - Higher Costs
  - SW VMWare ESX & VMotion
    - Enterprise license helps here
  - Greater complexity (Affinity rules, VIC Access, etc)
  - Requires better coordination & control around shared physical resources
  - Lack of support for “esoteric” HW interfaces

Comparisons & Contrasts

- PI: Basic approach is redundancy
  - Multiple instances of PI (server and/or interface running) but “bound” to HW
- VM: Separation from physical HW
  - Single instance of application – But “divorced” from physical HW
- Our Approach – **Use both**, to get more “flexibility” around physical resources, and redundancy of multiple instances of PI
- Tradeoffs: Cost and Complexity

OSIsoft®

Final Architecture – CLS design

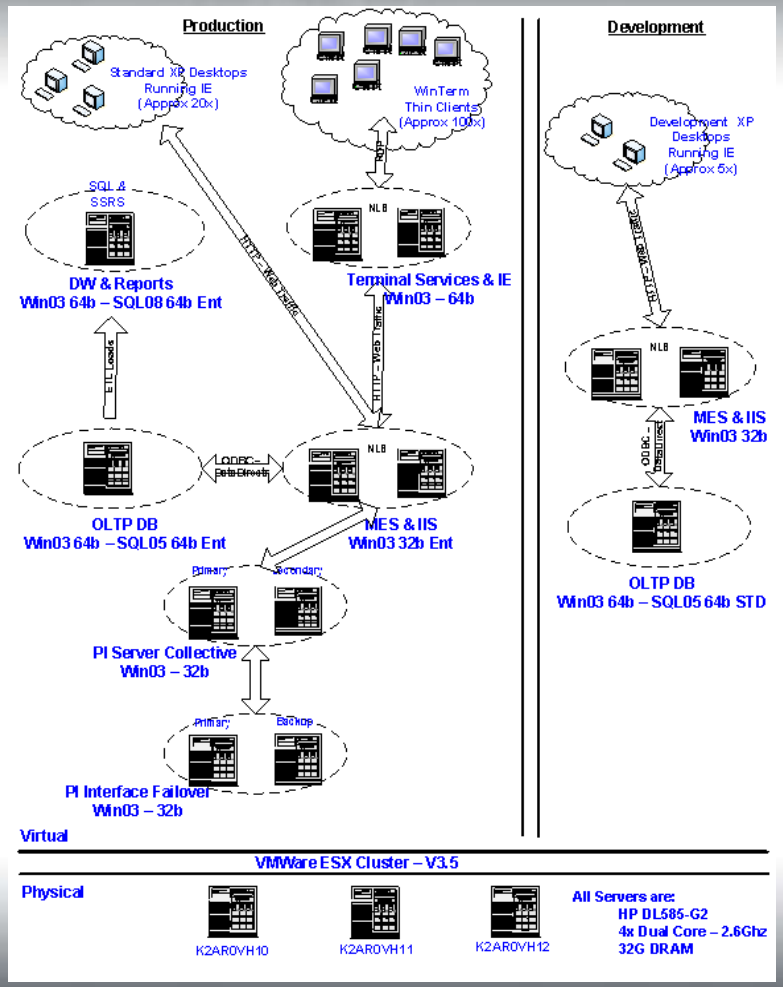

- 3 Physical Servers (all in ESX cluster)
- 18 “Guests”
  - 4 IT Servers
  - 3 Development Servers
  - 4 PI servers
  - 6 MES servers
  - 1 Other
- Designed to support shutdown of 1 physical server with no impact
- Use of VMware “affinity” rules

OSIsoft®

Final Architecture - Watchouts

- Need to define VMWare affinity rules - To keep redundant PI nodes on separate physical hardware in ESX cluster

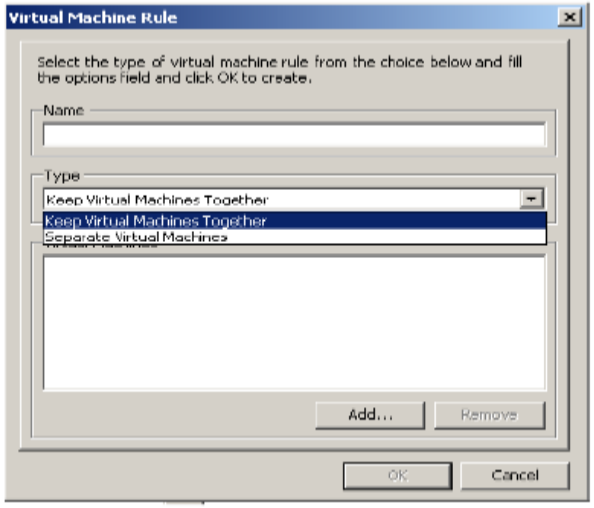

- We also use Network Load Balancing (NLB) on other server pairs, for other failover configurations – Significant setup complexity with server & network teams around physical & virtual VLAN configuration

Usage to date & Results

- Been operational for approximately 2 years
  - One unplanned PI interface outage (about 20 mins) due to failure around misconfiguration of shared interface file (≈ 99.999% uptime)
- VMWare and VMotion (all performed “while hot”)
  - Added third physical node to ESX cluster for scalability
  - Numerous virtual server moves (“VMigrations”) for load balancing
  - VMigration to move active PI servers to replace failing backup HW
  - Have had NO downtime associated with VMWare or VMotion
- PI HA
  - Migrated interface nodes from XP to virtual Win03 servers while “hot”
  - Support for numerous server and interface node reboots for applying monthly windows security patches & OPC interface upgrades
  - “Unplanned” power outage corrupted PI databases on PI HA Primary Was able to “promote” Sec->Pri, and initialize “old” Pri – with no outage

Future Needs and Plans

- Corning needs from OSI
  - Continue to make more interfaces support failover
  - N-way fanout of data from custom application writes via SDK rather than just API
  - Server-to-server replication – Need to address limits of HA V1
- Corning Plans
  - Testing new ACE beta, desire for faster failover
  - Bring up test AF server to test failover architectures (SQL server now in the mix) and possible replacement for batch DB

Future / Next Steps

- VMWare Futures (outside of PI HA)
  - VSphere – V4 – Fault Tolerance (FT) feature
    - Sets up “shadow” VM, kept in lockstep, transparent failover

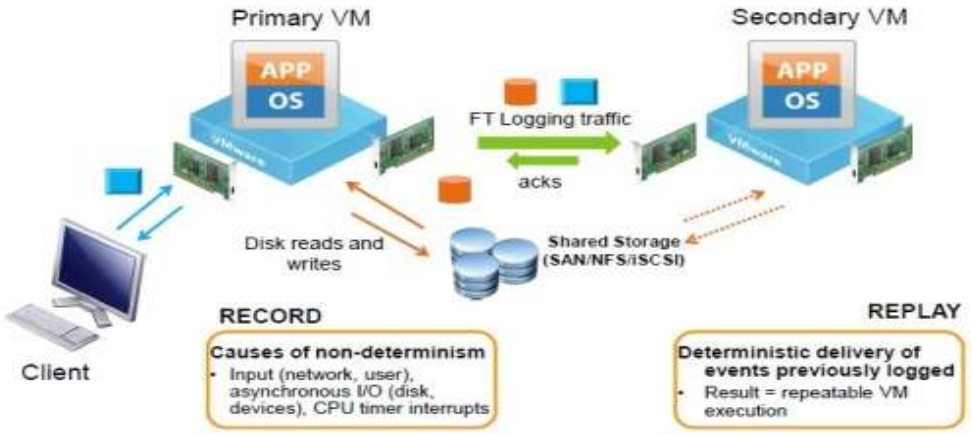

    - Limits
      - Single CPU – No SMP support (yet)
      - Survive HW crash, currently painful to do patching/upgrades

Thank You!

OSIsoft®