

PI System Tuning and Optimization

Presented By

Jay Lakumb, Product Manager
jlakumb@osisoft.com

Denis Vacher, Development Lead
dvacher@osisoft.com



OSIsoft®

Agenda

- Types of PI Systems
 - Hardware, Virtualization, and Critical Resources
- PI Archive Performance
 - Exception & Compression: Myths and Mistakes
 - Archive Sizing and Reprocessing
 - Demo
- Performance Monitoring Tips
- Real World Scenarios
- Optimization Techniques
 - General, Interfaces, Servers, Data Access, Clients

System Size... “YMMV”

- Point and Asset Count (S, M, L, XL, XXL)
 - 10K, 50K, 250K, 750K, 1M
- Data Rates
 - Common R:W ratio
 - Heavy read workload
 - Write-only systems
- Concurrent Users
 - Ranging from 1 to 1000
 - Dashboards/reports
- Server Applications
 - Batch, Performance Equations, Totalizer/Alarm
 - ACE, Notifications, MCN Health/IT Monitor
 - PI-to-PI Replication




rt
AMS/Phasors

alytics/BI

OSIsoft®

PI Server Hardware Allocation



OSIsoft®

Hardware Recommendations

- 64-bit platform, Windows 2008 R2 recommended
- Memory: 2GB minimum, up to 8KB per point
- CPU: as many cores as possible, based on client load
- Network: dedicated NIC(s), latency is most important
- Disk Controllers/Interface: key factors are I/O capacity and latency over bandwidth
- Storage: could be multi-tier, est. 5-20 bytes per event
- HA: PI Collectives mix up well with virtualization*

(*) assuming different hardware hosts per collective member



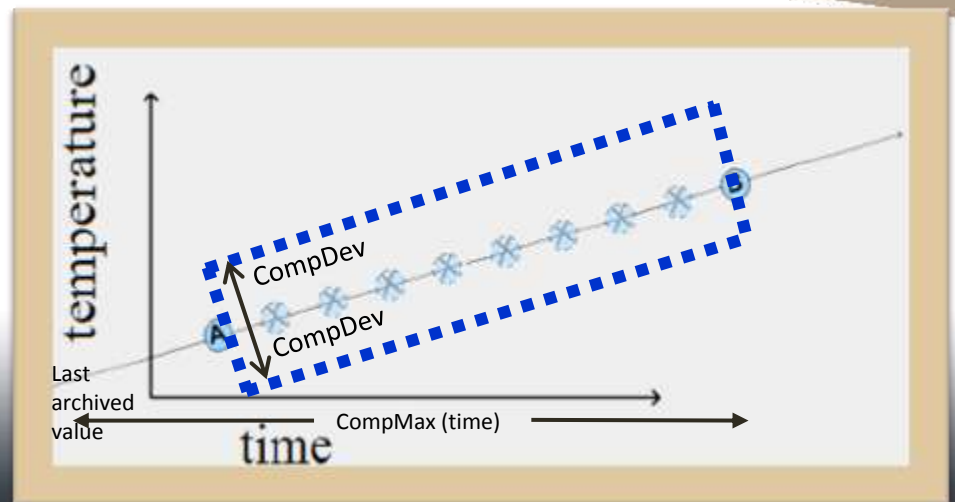
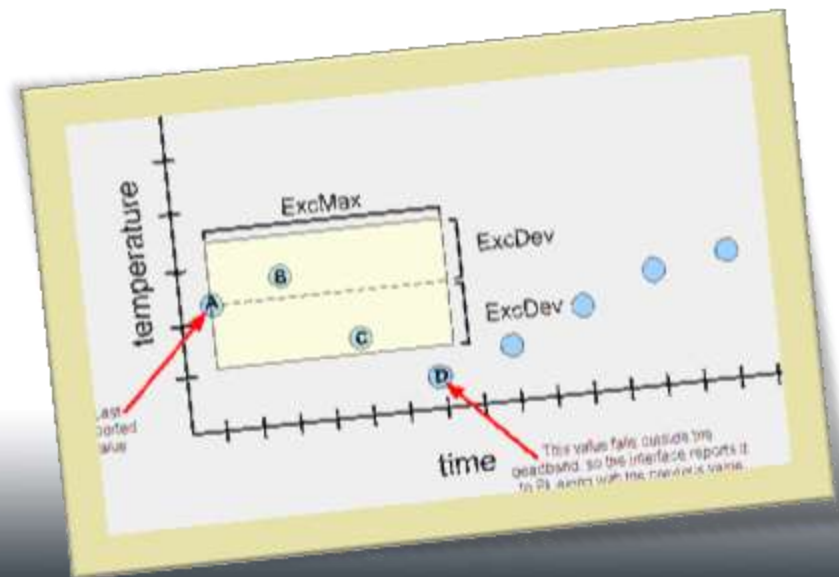

PI Server Virtualization

- 64-bit platform, Windows 2008 R2 recommended
- Memory: 2GB minimum, up to 8KB per point
- CPU: as many cores as possible, based on client load
- Network: dedicated NIC(s), latency is most important
- Disk Controllers/Interface: key factors are I/O capacity and latency over bandwidth
- Storage: could be multi-tier, est. 5-20 bytes per event
- HA: PI Collectives mix up well with virtualization*

(*) assuming different hardware hosts per collective member




PI Archive Performance Tuning



Efficient Signal Processing

- ❶ Raw Values
- ❷ After Exception
- ❸ After Compression



Exception & Compression

- Exception Processing (PI Interfaces)
 1. Eliminates Instrument Noise
 2. Conserves Network Bandwidth
 3. Conserves Server Resources
- Real-time Compression (Buffer/Snapshot Subsystems)
 1. Eliminates Process Noise
 2. Conserves Disk & I/O Resources
 3. Increases Overall System Performance
- High Fidelity & Efficiency
 - ☑ No Signal Averaging or Approximation
 - ☑ Preservation of all Minima/Maxima
 - ☑ No Additional Signal Latency



Demo



- Benefits of Exception/
Compression Tuning

OSIsoft®

Basic PI Server Tuning

- **Security – Auditing**

- Authentication Protocols `Server_AuthenticationPolicy`
- Administrative PI Trusts `AutoTrustConfig` (recommended: 0)
- Audit Trail `EnableAudit`

- **Archive/Queue Management**

- Auto Archive Creation `Archive_AutoArchiveFileRoot`
- Event Queue Settings `Snapshot_EventQueuePath`
 `Snapshot_EventQueueSize` (MB)

- **Network Connections**

- Stale Connection `MaxConnIdleTime` (seconds)

- **Updates**

- Don't Change These `MaxUpdateQueue`
 `TotalUpdateQueue`



Archive File Sizing

- Fixed Size or Dynamic?
 - Fixed size, with “auto-dynamic” growth
(Note: Fixed size required for automatic archive creation.)
 - Keep dynamic archives for backfilling jobs only
 - See [KB Article # 2998OSI8](#)
- Sizing Rule of Thumb
 - Size in MB $\geq 4,000 \times$ (Point Count)
 - For instance: 50,000 points \Rightarrow 200MB
- Optimal Sizing
 - No points with more than ~10 index records



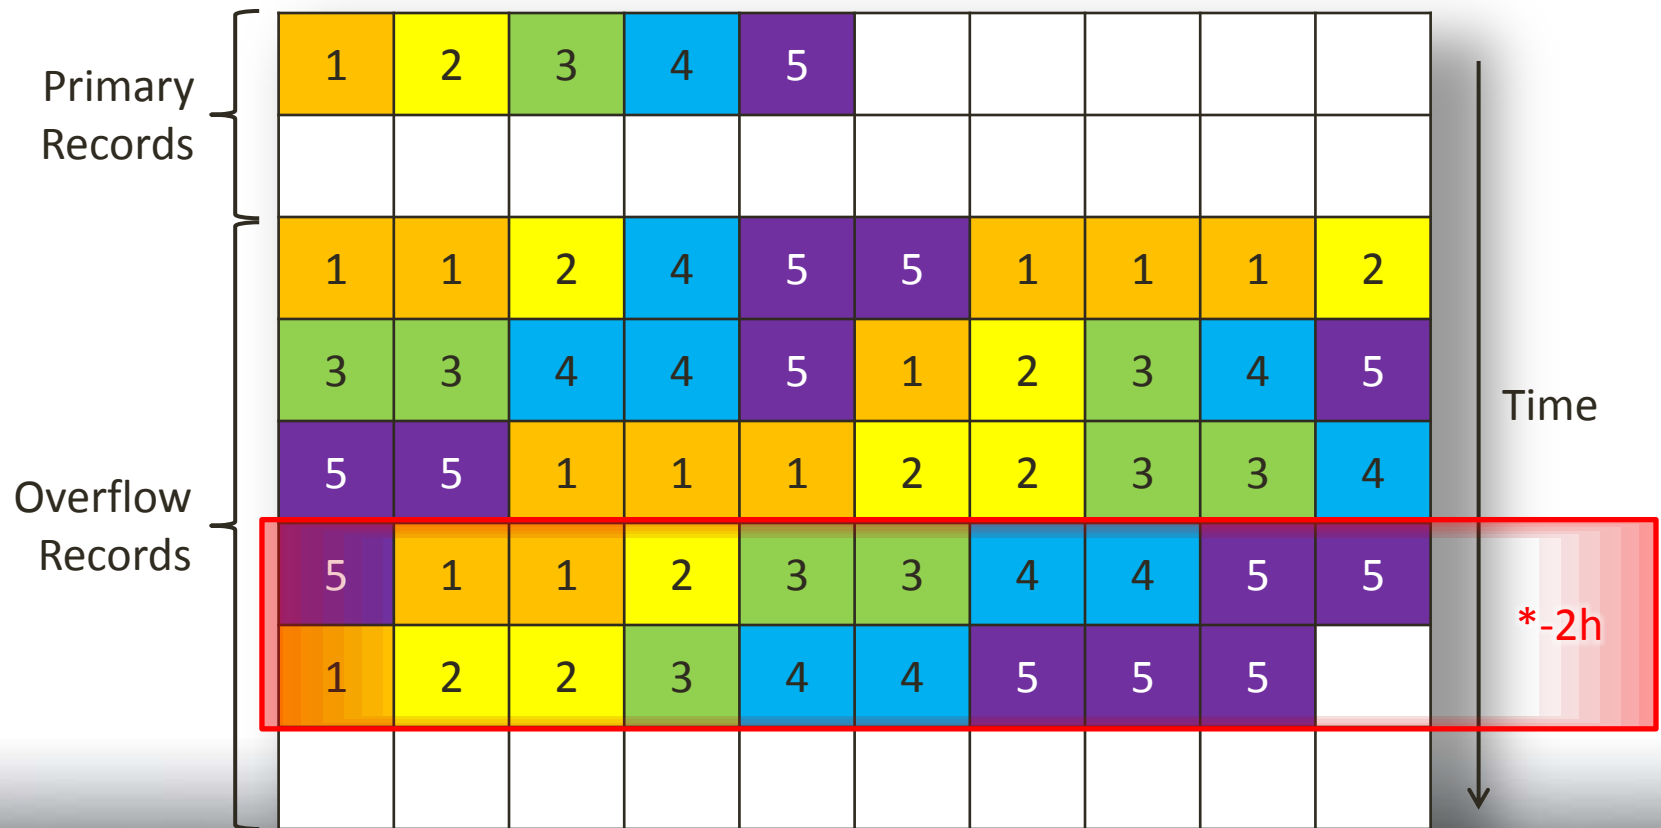
OSIsoft®

Archive Storage Map

```
Untitled - Notepad
File Edit Format View Help
D:\PI\arc\backup>c:\pdiag.exe -archk 2007.4GB
Analyzing archive: 2007.4GB
-----
recno: 1 id
recno: 2 id
recno: 3 id
recno: 4 id
recno: 5 id
recno: 6 id
recno: 7 id
recno: 8 id
recno: 9 id
recno: 10 id
recno: 11 id
recno: 12 id
recno: 13 id
recno: 14 id
recno: 15 id
recno: 16 id
recno: 17 id
recno: 18 id
recno: 19 id
recno: 20 id
recno: 21 id
recno: 22 id
recno: 23 id


Untitled - Notepad
File Edit Format View Help
D:\PI\arc\backup>c:\pdiag.exe -archk 2006.4GB
Analyzing archive: 2006.4GB
-----
recno: 1 id: 1 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 2 id: 2 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 3 id: 3 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 4 id: 4 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 5 id: 5 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 6 id: 6 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 7 id: 7 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 8 id: 8 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 9 id: 9 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 10 id: 10 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 11 id: 11 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 12 id: 12 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 13 id: 13 indices: 0 records: 1 events: 0 fr: 0.0%
recno: 14 id: 14 indices: 13 records: 1741 events: 326811 fr: 99.6%
recno: 15 id: 15 indices: 13 records: 1735 events: 325396 fr: 99.6%
recno: 16 id: 16 indices: 13 records: 1734 events: 325307 fr: 99.6%
recno: 17 id: 17 indices: 13 records: 1738 events: 326028 fr: 99.6%
recno: 18 id: 18 indices: 13 records: 1742 events: 326730 fr: 99.6%
recno: 19 id: 19 indices: 13 records: 1744 events: 327134 fr: 99.6%
recno: 20 id: 20 indices: 13 records: 1739 events: 325947 fr: 99.5%
recno: 21 id: 21 indices: 13 records: 1741 events: 326590 fr: 99.6%
recno: 22 id: 22 indices: 13 records: 1737 events: 325590 fr: 99.5%
recno: 23 id: 23 indices: 13 records: 1742 events: 326816 fr: 99.6%
```

Archive File Contents




OSIsoft®

File System Cache Efficiency?



After Archive Reprocessing



PI Server “KPIs”

Performance Counters - PI System Management Tools

File View Tools Help

Collectives and Servers

Search

Servers

- ☐ VCL10A

System Management Tools

Search

- Alarms
- Batch
- Data
- Interfaces
- IT Points**
 - Performance Counters**
 - Ring
 - SNMP
- Operation
- Points
- Security

Tag Settings Build Tags

Template Management

Load Apply Create UnLoad

Current Template(s):

Apply Counters to Instances

PI Archive Subsystem

- ☐ ArcEvent Calls/sec
- ☐ Archive Cycles/Sec
- ☒ Archived Events/sec
- ☐ Cache Clean Operations/sec
- ☐ Cache Flush Operations/sec
- ☒ Cache Record Count
- ☐ Cache Record Disk Reads/sec
- ☐ Cache Record Memory Reads/sec
- ☐ Cache Records Created/sec

Rate of archive events read.

Build?	Point Name	Performance Counter
<input checked="" type="checkbox"/>	VCL10A_PI Archive Subsystem_Cache Record Count	\\VCL10A\PI Archive Sub
<input checked="" type="checkbox"/>	VCL10A_PI Archive Subsystem_Archived Events/sec	\\VCL10A\PI Archive Sub
<input checked="" type="checkbox"/>	VCL10A_PI Archive Subsystem_Read Operations/s...	\\VCL10A\PI Archive Sub
<input checked="" type="checkbox"/>	VCL10A_PI Archive Subsystem_Events Read/sec	\\VCL10A\PI Archive Sub

Tag Count: 4

Build Tags

☒ Create tags on PI Server:

☐ Write tags to CSV File

Create Tags

OSIsoft®

K²PIs

Windows Counters

- **Core Process Health**
 - Process(...)\Creating Process ID
 - Process(...)\%Processor Time
 - Process(...)\Private Bytes
 - Process(...)\Working Set
 - Process(...)\Page Faults/sec
 - Thread(...)\ID Process
 - Thread(...)\Context Switches/sec
- **Overall CPU**
 - Processor\% Processor Time
- **Overall Memory**
 - Memory\Pages/sec
 - Memory\%Committed Bytes In Use
 - Memory\System Cache Resident Bytes
- **Overall Disk**
 - PhysicalDisk(...)\Avg. Disk Queue Length
 - PhysicalDisk(...)\Disk Transfers/sec
- **Overall Network**
 - Network Interface(...)\Bytes Total/sec
 - Network Interface(...)\Packets Received Errors

PI Server Counters

- **General Health**
 - PI Server Statistics(...)\Is Available
 - PI Server Statistics(...)\Is In Sync
 - PI Network Manager(_Total)\Connections
 - PI Network Manager(...)\Bytes Received/sec
 - PI Network Manager(...)\Bytes Sent/sec
- **Core Subsystem Health**
 - PI Subsystem Statistics(...)\RPC Request in Queue
 - PI Subsystem Statistics(...)\RPC Thread Active
 - PI Subsystem Statistics(...)\Transaction Completed/sec
 - PI Session Statistics(...)\Messages Received/sec
 - PI Session Statistics(...)\Messages Sent/sec
- **Data Flow Monitoring**
 - PI Snapshot Subsystem\Snapshots/sec
 - PI Snapshot Subsystem\OutOfOrderSnapshots/sec
 - PI Snapshot Subsystem\Queued Events/sec
 - PI Snapshot Subsystem\Events in Overflow Queues
 - PI Archive Subsystem\Archived Events/sec
 - PI Archive Subsystem\Total Unflushed Events
 - PI Archive Subsystem\Cache Flush Operations/sec
 - PI Archive Subsystem\Events Read/sec
 - PI Archive Subsystem\Current Cache Record Pool

PI Archive Tuning

PI System Management Tools

File View Tools Help

Collectives and Servers

Search

Servers

☒ VCL10A

System Management Tools

Search

Network Manager St
PI Services
PI Version
Reason Tree
Snapshot and Archiv
Tuning Parameters
Points
Digital States
Performance Equatio
Point Builder
Point Classes
Point Source Table
Totalizers

General Archive Backup Base Net Manager Snapshot Update Manager Security

Server	Collective	Name	Value	Default	Min	Max
VCL10A		Archive_PointLockTimeout		120,000	1,000	270,000
VCL10A		Archive_PointLockLogging		5,000	1	900,000
VCL10A		Archive_FlushThreadCount		4	1	128
VCL10A		Archive_FlushQueueSize		256	16	8,192
VCL10A		Archive_DataCoercionPolicy		0	0	3
VCL10A		Archive_LogArchivePercentFull		10	0	50
VCL10A		Plarchss ThreadCount		8	1	255
VCL10A		Archive CacheRecordPool		Point Count	10,000	500,000
VCL10A		Archive_MaxAnnotations		65,535	128	134,217,728
VCL10A		Archive_MaxWriteCachePerPoint		256	0	16,384
VCL10A		Archive_MinMemAvail		96	0	1,024
VCL10A		Archive_MinPercentMemAvail		10	0	50
VCL10A		Archive_SecondsBetweenFlush		0	1	86,400
VCL10A		Archive_ShirtFreeTime		1,600	0	86,400
VCL10A		Archive_ShortFreeTime		50	4	2,048
VCL10A		ArcMaxCollect		150,000	2,048	2,000,000,000
VCL10A		MarkArchiveGaps		1	0	1
VCL10A		Plarchss_QueuedMessageTimeout		600	10	3,600
VCL10A		Archive_AutoArchiveFileExt		.arc	0	0
VCL10A		Archive_AutoArchiveFileFormat		1	0	2
VCL10A		Plarchss	10,000	100,000	1	10,000,000

Size of the subsystem message thread pool. Message threads are dedicated to RPC request processing. Depending on the number of processors on the machine, this value may be increased so more RPC requests can be handled simultaneously. If all the threads are busy, RPCs are queued up and processed in chronological order. The subsystem limits COM Connector RPC request processing threads (called pending CC threads) to maximum of half of this number.

Plarchss_ThreadCount

Windows Counters

- **Core Process Health**

Process(...)\Creating Process ID
Process(...)\%Processor Time
Process(...)\Private Bytes
Process(...)\Working Set
Process(...)\Page Faults/sec
Thread(...)\ID Process
Thread(...)\Context Switches/sec

- **Overall CPU**

Processor\% Processor Time

- **Overall Memory**

Memory\Pages/sec
Memory\%Committed Bytes In Use
Memory\System Cache Resident Bytes

- **Overall Disk**

PhysicalDisk(...)\Avg. Disk Queue Length
PhysicalDisk(...)\Disk Transfers/sec

- **Overall Network**

Network Interface(...)\Bytes Total/sec
Network Interface(...)\Packets Received Errors

PI Server Counters

- **General Health**

PI Server Statistics(...)\Is Available
PI Server Statistics(...)\Is In Sync
PI Network Manager(_Total)\Connections
PI Network Manager(...)\Bytes Received/sec
PI Network Manager(...)\Bytes Sent/sec

- **Core Subsystem Health**

PI Subsystem Statistics(...)\RPC Request in Queue
PI Subsystem Statistics(...)\RPC Thread Active
PI Subsystem Statistics(...)\Transaction Completed/sec
PI Session Statistics(...)\Messages Received/sec
PI Session Statistics(...)\Messages Sent/sec

- **Data Flow Monitoring**

PI Snapshot Subsystem\Snapshots/sec
PI Snapshot Subsystem\OutOfOrderSnapshots/sec
PI Snapshot Subsystem\Queued Events/sec
PI Snapshot Subsystem\Events in Overflow Queues
PI Archive Subsystem\Archived Events/sec
PI Archive Subsystem\Total Unflushed Events
PI Archive Subsystem\Cache Flush Operations/sec
PI Archive Subsystem\Events Read/sec
PI Archive Subsystem\Current Cache Record Pool

Archive_CacheRecordPool

Windows Counters

- **Core Process Health**

Process(...)\Creating Process ID

Process(...)\%Processor Time

Process(...)\Private Bytes

Process(...)\Working Set

Process(...)\Page Faults/sec

Thread(...)\ID Process

Thread(...)\Context Switches/sec

- **Overall CPU**

Processor\% Processor Time

- **Overall Memory**

Memory\Pages/sec

Memory\%Committed Bytes In Use

Memory\System Cache Resident Bytes

- **Overall Disk**

PhysicalDisk(...)\Avg. Disk Queue Length

PhysicalDisk(...)\Disk Transfers/sec

- **Overall Network**

Network Interface(...)\Bytes Total/sec

Network Interface(...)\Packets Received Errors

PI Server Counters

- **General Health**

PI Server Statistics(...)\Is Available

PI Server Statistics(...)\Is In Sync

PI Network Manager(_Total)\Connections

PI Network Manager(...)\Bytes Received/sec

PI Network Manager(...)\Bytes Sent/sec

- **Core Subsystem Health**

PI Subsystem Statistics(...)\RPC Request in Queue

PI Subsystem Statistics(...)\RPC Thread Active

PI Subsystem Statistics(...)\Transaction Completed/sec

PI Session Statistics(...)\Messages Received/sec

PI Session Statistics(...)\Messages Sent/sec

- **Data Flow Monitoring**

PI Snapshot Subsystem\Snapshots/sec

PI Snapshot Subsystem\OutOfOrderSnapshots/sec

PI Snapshot Subsystem\Queued Events/sec

PI Snapshot Subsystem\Events in Overflow Queues

PI Archive Subsystem\Archived Events/sec

PI Archive Subsystem\Total Unflushed Events

PI Archive Subsystem\Cache Flush Operations/sec

PI Archive Subsystem\Events Read/sec

PI Archive Subsystem\Current Cache Record Pool

Archive_SecondsBetweenFlush

Windows Counters

- **Core Process Health**

Process(...)\Creating Process ID
Process(...)\%Processor Time
Process(...)\Private Bytes
Process(...)\Working Set
Process(...)\Page Faults/sec
Thread(...)\ID Process
Thread(...)\Context Switches/sec

- **Overall CPU**

Processor\% Processor Time

- **Overall Memory**

Memory\Pages/sec
Memory\%Committed Bytes In Use
Memory\System Cache Resident Bytes

- **Overall Disk**

PhysicalDisk(...)\Avg. Disk Queue Length
PhysicalDisk(...)\Disk Transfers/sec

- **Overall Network**

Network Interface(...)\Bytes Total/sec
Network Interface(...)\Packets Received Errors

PI Server Counters

- **General Health**

PI Server Statistics(...)\Is Available
PI Server Statistics(...)\Is In Sync
PI Network Manager(_Total)\Connections
PI Network Manager(...)\Bytes Received/sec
PI Network Manager(...)\Bytes Sent/sec

- **Core Subsystem Health**

PI Subsystem Statistics(...)\RPC Request in Queue
PI Subsystem Statistics(...)\RPC Thread Active
PI Subsystem Statistics(...)\Transaction Completed/sec
PI Session Statistics(...)\Messages Received/sec
PI Session Statistics(...)\Messages Sent/sec

- **Data Flow Monitoring**

PI Snapshot Subsystem\Snapshots/sec
PI Snapshot Subsystem\OutOfOrderSnapshots/sec
PI Snapshot Subsystem\Queued Events/sec
PI Snapshot Subsystem\Events in Overflow Queues
PI Archive Subsystem\Archived Events/sec
PI Archive Subsystem\Total Unflushed Events
PI Archive Subsystem\Cache Flush Operations/sec
PI Archive Subsystem\Events Read/sec
PI Archive Subsystem\Current Cache Record Pool

ArcMaxCollect

Windows Counters

- **Core Process Health**

Process(...)\Creating Process ID
Process(...)\%Processor Time
Process(...)\Private Bytes
Process(...)\Working Set
Process(...)\Page Faults/sec
Thread(...)\ID Process
Thread(...)\Context Switches/sec

- **Overall CPU**

Processor\% Processor Time

- **Overall Memory**

Memory\Pages/sec
Memory\%Committed Bytes In Use
Memory\System Cache Resident Bytes

- **Overall Disk**

PhysicalDisk(...)\Avg. Disk Queue Length
PhysicalDisk(...)\Disk Transfers/sec

- **Overall Network**

Network Interface(...)\Bytes Total/sec
Network Interface(...)\Packets Received Errors

PI Server Counters

- **General Health**

PI Server Statistics(...)\Is Available
PI Server Statistics(...)\Is In Sync
PI Network Manager(_Total)\Connections
PI Network Manager(...)\Bytes Received/sec
PI Network Manager(...)\Bytes Sent/sec

- **Core Subsystem Health**

PI Subsystem Statistics(...)\RPC Request in Queue
PI Subsystem Statistics(...)\RPC Thread Active
PI Subsystem Statistics(...)\Transaction Completed/sec
PI Session Statistics(...)\Messages Received/sec
PI Session Statistics(...)\Messages Sent/sec

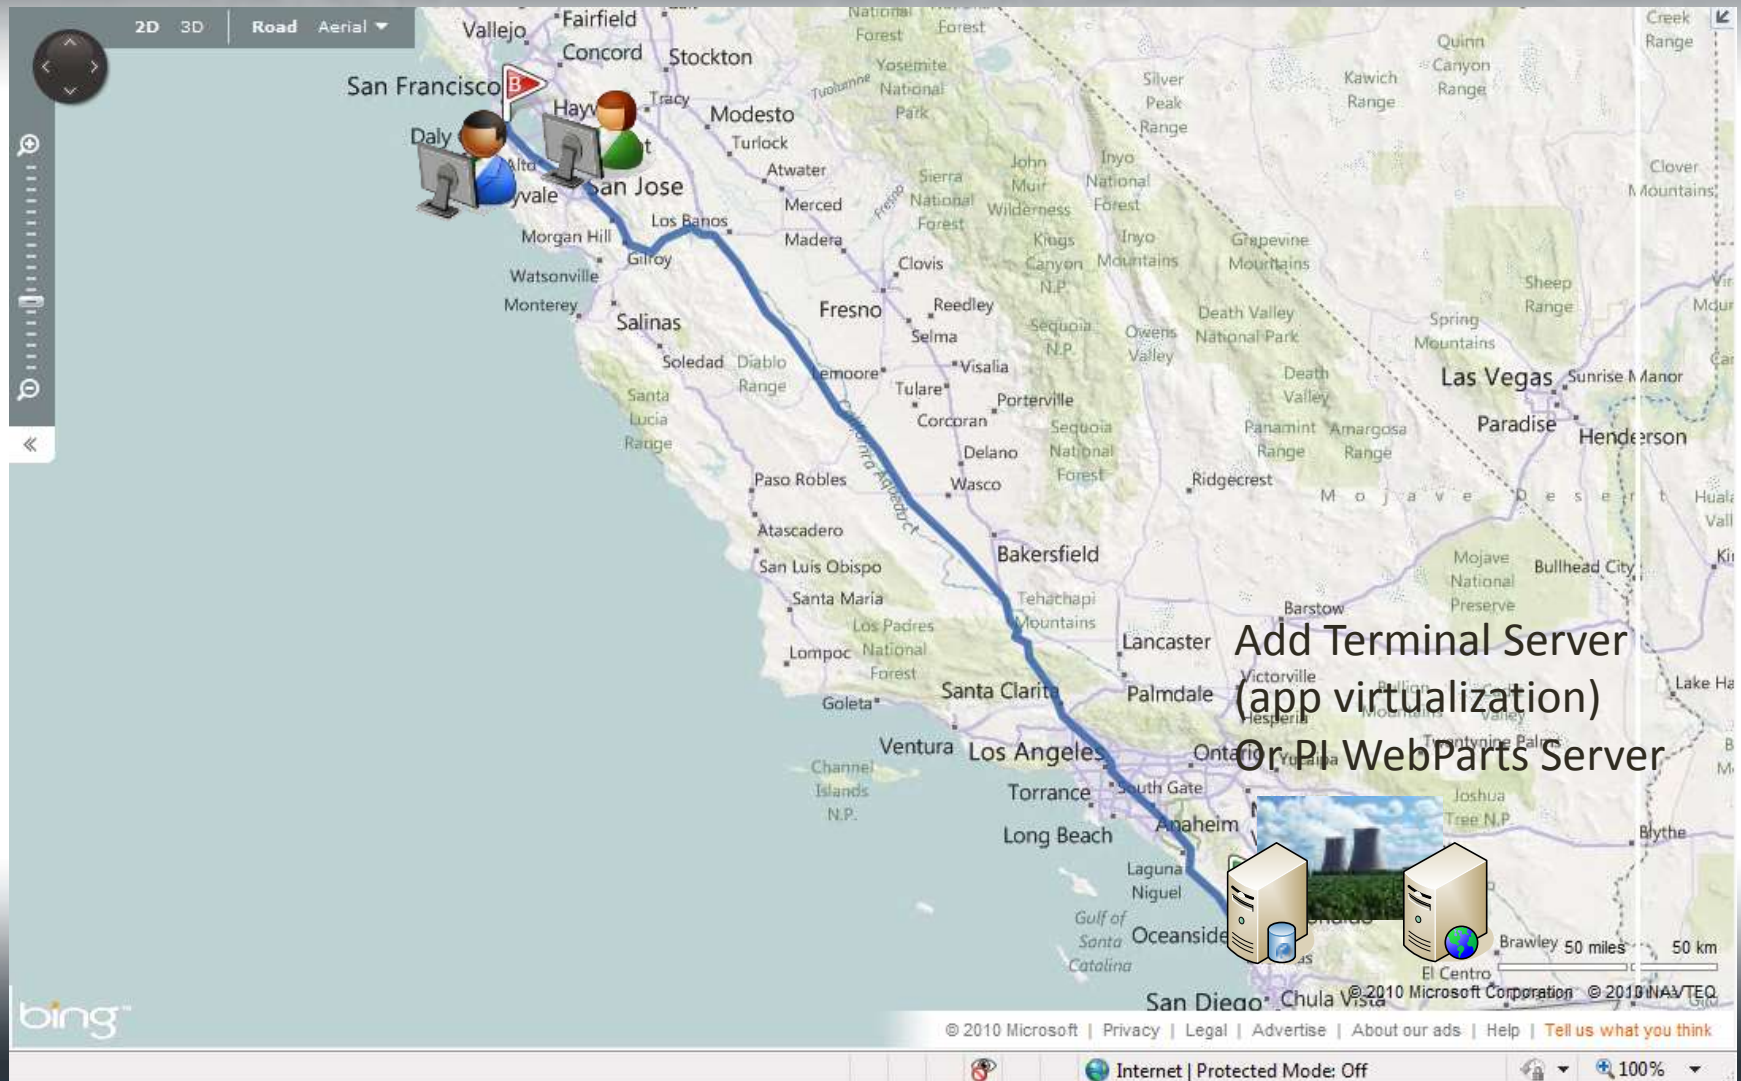
- **Data Flow Monitoring**

PI Snapshot Subsystem\Snapshots/sec
PI Snapshot Subsystem\OutOfOrderSnapshots/sec
PI Snapshot Subsystem\Queued Events/sec
PI Snapshot Subsystem\Events in Overflow Queues
PI Archive Subsystem\Archived Events/sec
PI Archive Subsystem\Total Unflushed Events
PI Archive Subsystem\Cache Flush Operations/sec
PI Archive Subsystem\Events Read/sec
PI Archive Subsystem\Current Cache Record Pool

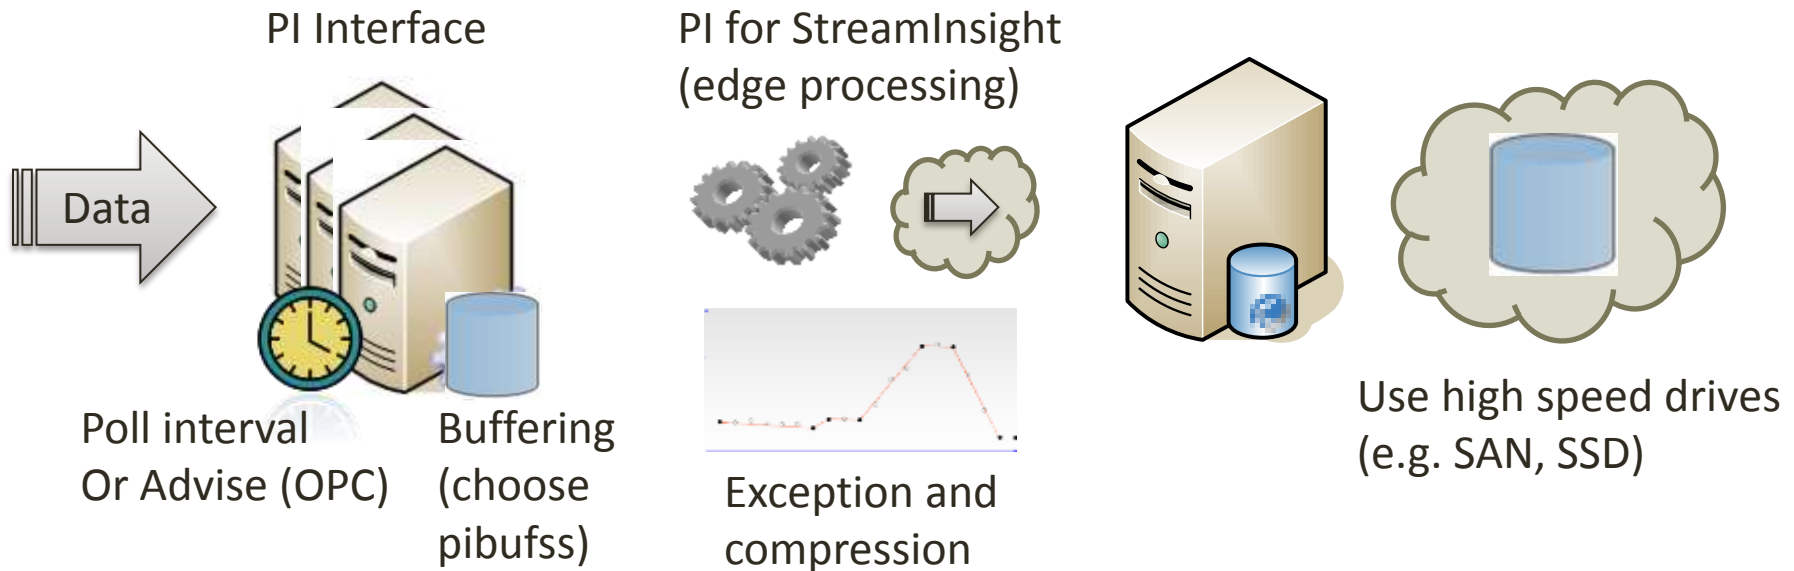
Satellite or WAN Connection




Geographical Distribution






Collect Large Volumes of Data




Access Large Volumes of Data




High speed drives
(e.g. SAN, SSD)



HA PI Secondary
(load distribution)




Scale out analytics
and data access



PI SDK techniques
(e.g. PlotValues)



PI for StreamInsight
(edge processing)



Optimization Techniques: General

- Ensure system/virtual machine is sized appropriately
 - “Hardware and System Sizing Recommendations” spreadsheet on Tech Support site
- Disable unused services, e.g. IIS, Indexing, etc. or use Windows Server Core
- Clamp or prevent expensive queries
- Provide sufficient network bandwidth
- Choose an optimal deployment configuration
- Upgrade to latest releases of PI and Microsoft products
- Leverage High Availability (PI/AF/Notif, Interface failover, SharePoint server farms, etc.) to achieve higher uptime
- Tools to identify performance bottlenecks: MCN Health Monitor, PI/SDK Message Log, Visual Studio Profiler



Optimization Techniques: Interfaces

- KB #00266 – “When using the PI to PI interface, is it better to push or pull the data?”
- Scale out (multiple instances) to handle more/faster data
- Use Buffering to minimize data loss
- Configure disconnected startup
- Choose appropriate scan times
- Separate instances for different scan times, output points
- Check for Stale or Bad points
- KB #00260 – “How does the number of Advise tags per group in OPC interface affect throughput?”



Optimization Techniques: Servers

- “PI Server Performance Monitoring” White Paper on TS site
- KB #00159 – “What is the Offline Archive Utility (piarchss) and how do you use it to reprocess archives?”
- Apply multiple conditions to each notification (not 1:1)
- Increase periodic time rule (ACE, PI Notifications), or try natural scheduling (uses snapshot caching)
- Minimize number of COM Connector tags
- Aggregate data (using PE or Totalizer tags)
- Scale out (PI, AF, PE, ACE, Notifications) to handle more users/queries with High Availability functionality
- Separate heavy and casual users with HA functionality



Optimization Techniques: Data Access

- “SQL Optimization” White Paper on TS site
- Use PI OLEDB and PI OLEDB Enterprise as in-process
- Break queries into smaller chunks (e.g. fewer tags/assets, shorter time ranges)
- Builder’s Café Webinar on “Optimizing your PI SDK apps”
- Use bulk queries (e.g. PI SDK ListData functions, AF SDK elements), PlotValues, asynchronous calls, multithreading
- Choose appropriate binding for PI Web Services
- Leverage PI Data Access products that implement (and abstract) advanced techniques such as connection pooling, caching, and event pipes (e.g. PI Web Services)



Optimization Techniques: Clients

- Move PI Calc/ODBC datasets to server (PE, AF, ACE)
- Minimize data loaded/displayed on startup – use drill down for more details
- Split large amount of data across multiple displays/pages
- Try HTTP Compression with PI WebParts
- Use SharePoint server farms, load balancing with WP
- Increase update interval (PB, DL, WP)
- Run PI Clients as close to data sources as possible – leverage application virtualization



Thank You!



OSIsoft®