
Holiday/Special Event Preparedness for WCS Sites Top Ten Actions

Stacy Joines

Distinguished Engineer

ISSW WebSphere Performance Consulting



High-Volume “Gaming Strategy”



- High-volume Web sites have two modes:
 - Achieving targets
 - Down

- High-capacity support and versatility are at odds with each other
 - Constant change in applications, features, and content reduces capacity

- The greatest competitive advantage is often just being available
 - Delivering core services (browse and buy) consistently beats delivering nothing
 - Protecting these core services requires planning

- Holiday planning activities start early in the year and carry throughout

Action #1: High-Risk Activities Come First




- 1st and 2nd Quarters are the time to *execute/deploy* high-risk projects
 - Major upgrades (stack and OS versions)
 - Major development projects (significant new features, new sites, etc.)
 - Significant topology changes (new hardware platforms, multi-datacenter, etc.)
 - Plan for any upgrades needed to the test beds to reflect major upgrades

- 3rd Quarter provides opportunity for stabilizing early changes and to complete low-risk projects
 - Identify any follow-on work, additional capacity, etc. required to harden earlier major projects
 - As early work stabilizes, complete medium to low risk projects
 - Also the quarter to begin execution of the holiday readiness plan
 - Determine growth year over year
 - Begin testing process.

- 4th Quarter
 - Changes hopefully limited to content by late September/October timeframe
 - Regular Monitoring and data analysis to highlight problem spots as the date draws closer
 - Preparing holiday plans for staffing, etc.
 - Ready plans and development for upcoming 1st & 2nd quarter development work

Action #2: Observe the Site before the Event



- Observe site visitor usage patterns
 - Use access log analysis or consider using a metrics service

- Regular observation
 - Establishes trend information regarding usage patterns and resource consumption
 - Identifies other, smaller peak events

- Consider if shifts in usages patterns impact performance
 - Updates to test approach?
 - Additional hardware required to support more personalized content delivery?

- Review the logs on all systems
 - Look for recurring errors or warnings. The issues behind these messages might amplify under peak loading
 - Consider “log sweeping” if the farm is large

Action #2: Observe the Site before the Event (cont'd)



- Find a monitoring system and use it
 - Setting up a new monitoring system takes time
 - Installation often spans multiple systems and requires scheduling
 - Alert threshold requires set-up and burn-in
 - Recommended as a task for 1st or 2nd quarter to give sufficient time
 - Operations team also requires training to use the monitoring system effectively


- Advanced topic: Integrate mid-tier and remote systems monitoring
 - Correlates events at the mid-tier with events throughout the environment
 - Supports a proactive vs. reactive model to potential outages

Activity #3: Get the Performance Test Bed in Shape

- Is the performance test bed a “scale model” of production?
 - Software/hardware upgrades and farm growth impact production size

- Mimic production with the performance test bed
 - Similar hardware base (CPU, memory, OS, etc.)
 - Software at same levels
 - Proportions are correct relative to the production farm
 - HttpServer – WCS – Database

- Consider other systems as well
 - Did the site add a new peripheral system this year?
 - Search engine, LDAP server, etc.
 - Provide test editions of all major peripheral systems
 - Maintain proportions whenever possible
 - If test harnesses simulate some systems, certify these systems independently



Testbed Proportion Example


If Production looks like this...



Http Layer
8 cores



WCS Layer
20 cores



Database Layer
40 cores

- Maintain similar ratios in the performance test environment
- For example, a 25% test environment would be:
 - Http Layer: 2 cores
 - WCS Layer: 5 cores
 - Database Layer: 10 cores
- Too little resource in one layer creates an artificial bottleneck
- Too much resource in one layer does not reproduce production scale

Activity #4: Include the Business in the Plan



- Business drivers often dictate web site operations and deployment schedules
- Understanding business plans impacts testing and site prep for the Event
- Promotions
 - Special offers: Look for anything that might drive significant load onto the site
 - Promotion scheduling: Plan for promotions starting earlier than normal or encompassing a different timeframe
 - New promotional interactions: Tie-ins with the stores; linkages other businesses drive change in usage patterns
 - Build test scenarios to explore the performance of proposed promotions
- Features and capabilities
 - New features and functions planned for holiday
 - How will they be used and tested?
- Reducing risk
 - Major new features arriving just before holiday require mitigation
 - Mitigate through deployment by deploying to non-mainline hardware
 - Mitigate through availability by designing “circuit breakers” to disable function under load if it becomes problematic

Activity #5: Develop an Event Response Plan



- Identify critical operational elements
 - Critical components: Catalog servers, databases, OMS systems, etc.
 - “Tender” spots with known issues: Remote systems, remote providers, infrastructure, etc.

- Plan for the *eventuality* of an outage in these systems during peak times
 - What are the likely symptoms?
 - What is the overall impact to the site (outage, limited functionality, etc.)?
 - What are the recovery procedures?
 - Who owns the various steps of the procedures?

- Take the overall team through practice drills prior to the holidays
 - Establishes roles and resources early
 - Allows everyone to gain confidence with their roles prior to holiday

- Response plan often reveals areas requiring further hardening
 - Asynchronous vs. synchronous coding patterns to prevent stalls
 - Better management of interactions with non-critical systems to prevent interruptions
 - Hardening/increasing capacity of critical systems

Activity #6: Plan for Holiday/Event Staffing

- Name the team(s) managing the Web site through holiday
 - Identify key players from major web site areas
 - Operations, DBAs, development, network, etc.
 - Schedule any vendor or partner resource required
 - Contractors, remote systems providers, hosting, etc.
 - Holiday is vacation time
 - Identify resources early to ensure coverage
- Setup communications channels in advance
 - Distribute bridge numbers and setup checkpoint calls in advance
 - Identify mgt/executive team coverage and escalation channels early
- Determine coverage duration and locality
 - 24x7 sites may require multiple shifts of coverage
 - Many organizations setup an on-site “warm-room” for Black Friday weekend support

Activity #7: Develop a Performance Test Plan

- Functional testing is valuable and necessary however...
 - Functional testing is typically single user testing
 - Not a valid predictor of performance or correctness under load from many users
- Likewise “normal” day site behavior is not a predictor of peak day behavior
 - *Question: Why not just multiply the resource consumption on a normal day and project?*
 - *Answer: Linear scale cannot be assumed; it must be verified*
- Code and data can have problems that prevent them from scaling linearly
 - For example, a 2x growth in data could require n^2 increase in CPU to process
 - Likewise, a synchronized method could drive response times increasingly higher
 - Neither problem would present itself at low data or user volumes
- A representative performance test is the best approach to predict capacity



Activity #7: Develop a Performance Test Plan (con'td)


- What are the critical elements for test this year
 - New features, functions, upgrades, systems
 - Anything that did not work well last year
- Establish resources needed and coverage gaps
 - Scripters, tester, analysts, etc.
 - Sufficient test licenses and test systems
- Define the coverage area for the existing environment and test bed
 - Does the site require additional external or internal testing to cover all function or volumes?
- Define a data plan as well
 - Sufficient catalog data, user identities, promotions, etc. to populate a meaningful test
- Timing
 - Test with enough lead time to remedy issues uncovered
 - However, major site changes should stabilize prior to testing

Activity #8: Evaluate Cache Needs



- Caching significantly increases site capacity
- Cache value increases when:
 - Frequently visited pages are cacheable
 - Large portions of pages are cacheable
 - Cached elements are pushed close to the end user (even outside the data center)
 - The cache refreshes infrequently, reducing the CPU burden to rebuild cached elements
- Reconcile capacity needs with business rules
 - Overly aggressive rules result in frequent data loads and resulting cache updates
- Consider relaxing rules for peak events
 - Keep items cached longer/reduce load frequency, especially during peak hours
- Caching strategy often broken by new features and content
 - Surplus “average” capacity compensates for most of the year
 - Needs attention prior to holiday to make “peak” capacity

Caching Benefit



- Cache as near to the user as possible
 - Reduces CPUs and other resources engaged to deliver the request
 - Also reduces latency and improves response time
- Caches become more expensive and restrictive to refresh near the edge
 - Cache content consists of larger and typically more frequently accessed pages
 - External cache providers may limit refresh frequency



Activity #9: Consider Industry Trends

- Look for products and technology trends for insight into potential customer interest areas at holiday
- Lower revenue does not necessarily imply decreasing visits
 - Revenue stream is tied to sales and other promotions
 - Many Web sites experienced significant increases in traffic last year despite the poor economy
- Consumers are developing their own gaming strategies around holiday shopping
 - Consumers leveraging Gift Registry and saved Shopping Carts to snag the best deals
 - Reserve inventory prior to sales via carts
 - Identify desirable items for distant relatives via Registry
 - Revisiting the site frequently throughout sales days to apply promotions

Activity #10: Getting Help



- Capacity and performance planning
 - IBM Tech Sales is a good source to assist with initial capacity planning
 - ISSW service offerings to assist with integrating performance into the project lifecycle

- Performance testing
 - ISSW provides SME support for customer in-house organizations
 - GTS partners with ISSW to provide broader performance support

- HA/DR planning
 - GTS/GBS partners with ISSW and other organizations to develop HA/DR strategies
 - IBM GTS SO organization provides hosting services

Activity #10: Getting Help (cont'd)



- IBM Premium Support Options
 - Support options and personnel tailored to your environment
 - Assist with managing service levels across the environment

- Training
 - ISSW offers “Center of Excellence” training around performance
 - Also planning workshops and other offerings available

Summary



- Performance objectives require on-going focus to achieve
- Start with core functionality and harden supporting systems/infrastructure
- Plan with core functionality in mind
 - Risks are understood and mitigated
 - Rollback plans defined and practiced
- Integrate performance regression into ongoing updates and fixes
- Start early to prepare for peak events
 - Learn from the existing traffic to prepare tests
 - Practice peak day activities under load
 - Validate performance with testing in a scale environment
 - Manage change tightly
 - Consider business rule flexibility to reduce resource burden
- Be in the game when it counts the most...

Questions?

